arXiv:1408.2758v1 [cs.CR] 12 Aug 2014

Karlsruhe Institute of Technology

The Influence of Architectural Styles on Security, Using the Example of a Certification Authority

Study Thesis by

Michael Tänzer

At the Department of Informatics
Institute for Program Structures
and Data Organization (IPD)

Reviewer:	Prof. Dr. Ralf H. Reussner
Second reviewer:	Prof. Dr. Walter F. Tichy
Adviser:	M. Sc. Zoya Durdik
Second adviser:	Dipl.-Inform. Matthias Huber

Duration: 15. May 2013 – 17. July 2013

I declare that I have developed and written the enclosed thesis completely by myself, and have not used sources or means without declaration in the text.

Karlsruhe, 16. July 2013

.....

(**Michael Tänzer**)

Abstract

Often, security is considered in an advanced stage of the implementation of a system, rather than integrating it into the system design. This leads to less secure systems, as the security mechanisms are only applied as an afterthought and therefore do not integrate well with the rest of the design. Also, several statistics about discovered vulnerabilities in existing systems suggest, that most of the vulnerabilities of a system are not caused by errors in the cryptographic primitives, but in other parts of the implementation. So integrating security concerns early in the design process seems a promising approach for increasing the security of the resulting system.

This work evaluates how the choice of the architectural style affects the security of the resulting system. The evaluation is done on the example of an existing certification authority (CA). The requirements for the system are gathered and multiple designs according to different architectural styles are drafted and evaluated using a risk evaluation method. Then the evaluated designs are compared to find out whether there are significant differences.

Zusammenfassung

Häufig wird die Sicherheit eines Systems erst in einem fortgeschrittenen Entwicklungsstadium berücksichtigt, statt dies in den Entwurfsprozess zu integrieren. Dies führt zu unsichereren Systemen, da die Sicherheitsmechanismen erst im Nachhinein hinzugefügt werden und deshalb nicht gut in den bestehenden Systementwurf integriert sind. Außerdem lassen einige Statistiken über entdeckte Sicherheitslücken in existierenden Systemen vermuten, dass die meisten Fehler die zu Sicherheitslücken führen, nicht in den kryptographischen Primitiven zu finden sind, sondern in anderen Teilen der Implementierung. Deshalb scheint die Integration von Sicherheits-Aspekten früh im Entwurfsprozess vielversprechend um die Sicherheit des resultierenden Systems zu verbessern.

Diese Arbeit evaluiert, wie sich die Entscheidung für einen Architekturstil auf die Sicherheit des resultierenden Systems auswirkt. Die Bewertung wird am Beispiel einer existierenden Zertifizierungsstelle (CA) ausgeführt. Die Anforderungen an das System werden erhoben und es werden mehrere Entwürfe anhand verschiedener Architekturstile angefertigt und anhand einer Risiko-Bewertungsmethode evaluiert. Anschließend werden die evaluierten Entwürfe verglichen um festzustellen ob es signifikante Unterschiede gibt.

Contents

Abstract iii

Zusammenfassung v

1. Introduction 1

2. Foundations 3
2.1. Definitions . . . 3
2.2. CAcert . . . 3
2.2.1. Features of the Existing System . . . 4
2.2.2. Problems of the Existing System . . . 5
2.3. Method for Evaluation . . . 6
2.4. Related Work . . . 7

3. Approach 9
3.1. Requirement Analysis . . . 9
3.2. Evaluation of Architectural Styles . . . 10
3.3. Comparison of the Architectural Styles . . . 10

4. Architectural Design Alternatives 11
4.1. Layered Architecture . . . 11
4.1.1. Choice . . . 11
4.1.2. Design . . . 11
4.1.2.1. Front-End Layer . . . 11
4.1.2.2. Authentication and Authorisation Layer . . . 12
4.1.2.3. Operations Layer . . . 13
4.1.2.4. Data Abstraction Layer . . . 13
4.1.2.5. Back-End/Supporting Services Layer . . . 13
4.2. Pipes and Filters Architecture . . . 14
4.3. Service-Oriented Architecture . . . 14
4.3.1. Choice . . . 14
4.3.2. Design . . . 14
4.3.2.1. Front-End . . . 14
4.3.2.2. Authentication and Authorisation . . . 14
4.3.2.3. Business Logic . . . 15

5. Evaluation 17
5.1. Variation of the Evaluation Method . . . 17
5.2. Layered Architecture . . . 17
5.2.1. Security Evaluation . . . 17
5.2.2. Overall Evaluation . . . 18
5.3. Service-Oriented Architecture . . . 18
5.3.1. Security Evaluation . . . 18
5.3.2. Overall Evaluation . . . 19

5.4. Comparison . . . 19

6. Conclusion **21**

Bibliography **23**

Glossary **25**

A. Requirements **27**
A.1. Functional Requirements . . . 27
A.2. Non-Functional Requirements . . . 43
A.3. Critical Assets . . . 43
A.3.1. Web of Trust . . . 43
A.3.1.1. Integrity . . . 43
A.3.1.2. Confidentiality . . . 44
A.3.2. Login/Recovery Credentials . . . 44
A.3.2.1. Integrity . . . 44
A.3.2.2. Confidentiality . . . 44
A.3.3. User Data . . . 44
A.3.3.1. Integrity . . . 44
A.3.3.2. Confidentiality . . . 44
A.3.4. Issued Certificates . . . 44
A.3.4.1. Integrity . . . 44
A.3.4.2. Confidentiality . . . 45
A.3.5. Revocation Information . . . 45
A.3.5.1. Integrity . . . 45
A.3.5.2. Availability . . . 45
A.3.6. Root/Subroot Certificates . . . 45
A.3.6.1. Integrity . . . 45
A.3.7. Certificate Signing Keys . . . 45
A.3.7.1. Integrity . . . 45
A.3.7.2. Confidentiality . . . 45

B. Layered Design **47**

C. Layered Design Analysis **49**
C.1. Web of Trust . . . 49
C.1.1. Unauthorised Modification of the Web of Trust . . . 49
C.1.2. Violation of the Confidentiality of the Web of Trust . . . 51
C.2. Login Credentials . . . 52
C.2.1. Unauthorised Modification of the Login Credentials . . . 52
C.2.2. Violation of the Confidentiality of the Login Credentials . . . 52
C.3. User Data . . . 53
C.3.1. Unauthorized Modification of User Data . . . 53
C.3.2. Unauthorized Access to User Data . . . 53
C.4. Issued Certificates . . . 53
C.4.1. Modification of Issued Certificates . . . 53
C.4.2. Unauthorized Access to Issued Certificates . . . 53
C.5. Revocation Information . . . 54
C.5.1. Unauthorized Modification of Revocation Information . . . 54
C.5.2. Prevent Access to Revocation Information . . . 54
C.6. Root/Subroot Certificates . . . 55
C.6.1. Modification of Root/Subroot Certificates . . . 55

C.7. Certificate Signing Keys . . . 55
C.7.1. Modification of the Certificate Signing Keys . . . 55
C.7.2. Violation of the Confidentiality of the Certificate Signing Keys . . . 55

D. Service-Oriented Design **57**

E. Service-Oriented Design Analysis **59**
E.1. Web of Trust . . . 59
E.1.1. Unauthorised Modification of the Web of Trust . . . 59
E.1.2. Violation of the Confidentiality of the Web of Trust . . . 60
E.2. Login Credentials . . . 61
E.2.1. Unauthorised Modification of the Login Credentials . . . 61
E.2.2. Violation of the Confidentiality of the Login Credentials . . . 61
E.3. User Data . . . 62
E.3.1. Unauthorized Modification of User Data . . . 62
E.3.2. Unauthorized Access to User Data . . . 62
E.4. Issued Certificates . . . 62
E.4.1. Modification of Issued Certificates . . . 62
E.4.2. Unauthorized Access to Issued Certificates . . . 63
E.5. Revocation Information . . . 63
E.5.1. Unauthorized Modification of Revocation Information . . . 63
E.5.2. Prevent Access to Revocation Information . . . 63
E.6. Root/Subroot Certificates . . . 64
E.6.1. Modification of Root/Subroot Certificates . . . 64
E.7. Certificate Signing Keys . . . 64
E.7.1. Modification of the Certificate Signing Keys . . . 64
E.7.2. Violation of the Confidentiality of the Certificate Signing Keys . . . 64

1. Introduction

The lack of design with security in mind has often been discussed [FP01, And01, Sta11, p. 38]. As a solution, processes that take security into account in all aspects of the software development process and try to avoid this "security as an afterthought" problem have been proposed [MGM03]. Additionally, statistics of vulnerabilities found in existing software systems suggest that it is not so much deficiencies in using cryptographic primitives, but errors in the general implementation of the software that cause most of the vulnerabilities. According to [CCCN99, p. 110] less than 15% of all computer emergency response team (CERT) advisories could have been avoided by proper use of cryptography and statistics of the common vulnerabilities and exposures (CVE) index [CM07] indicate that only 1.5% of all vulnerabilities were caused by errors in parts of the software concerning cryptography. Those vulnerabilities may also partly stem from a design that has the tendency to be less secure when implemented.

These observations led to the consideration of evaluating different architectural styles with respect to the security likely to be achieved by an implementation of these styles. Because such an evaluation is pretty hard if not impossible to do on a "general idea" of the architectural style, a real system with real requirements (some of the main non-functional requirements relating to security of the system) is chosen and some designs according to different architectural styles are crafted and evaluated. The concrete system selected for this work is a CA based on a web of trust, named CAcert, which is described in Section 2.2. CAcert was chosen because of the high security requirements needed in this environment and the author's familiarity with that system.

This work focuses on the design on the architectural level and evaluating the resulting design, it does not go further into a detailed design or even an implementation stage. Due to this limitation of scope a formal validation of the approach is not conducted. Also due to time constraints only two architectural styles were evaluated, however this is not an inherent limit and the method provided allows to compare more styles with each other.

This work is structured as follows. Chapter 2 introduces basic concepts and gives an overview of related work. Chapter 3 describes the approach used in this work. In Chapter 4, several designs according to different architectural styles are drafted, which are then evaluated in Chapter 5. Finally in Chapter 6, the conclusions that can be drawn from the evaluation, are presented.

2. Foundations

This work is based on various concepts from other works which are introduced in this chapter. In Section 2.1 definitions about what constitutes an architectural style and basic security properties are established. Section 2.2 describes the CAcert system in more detail, from the basic principle of operation to a description of the existing software and its shortcomings. The attack tree method used for evaluating the designs is presented in Section 2.3, and Section 2.4 lists work related to the topic of this paper.

2.1. Definitions

Definition 1 (Architectural Style)**.** "If architecture is a formal arrangement of architectural elements, then architectural style is that which abstracts elements and formal aspects from various specific architectures. An architectural style is less constrained and less complete than a specific architecture. [...] The important thing about an architectural style is that it encapsulates important decisions about the architectural elements and emphasizes important constraints on the elements and their relationships. The useful thing about style is that we can use it both to constrain the architecture and to coordinate cooperating architects." [PW92]

Definition 2 (Confidentiality)**.** "The property that information is not disclosed to system entities (users, processes, devices) unless they have been authorized to access the information." [Com10]

Definition 3 (Integrity)**.** "The property whereby an entity has not been modified in an unauthorized manner." [Com10]

Definition 4 (Availability)**.** "The property of being accessible and useable upon demand by an authorized entity." [Com10]

2.2. CAcert

CAcert is a free[1] CA, driven by an open community. That means it offers X.509 and other certificates that users can use for various purposes including securing network traffic with the transport layer security (TLS) protocol.

CAcert operates on a non-commercial basis which is made possible by a web of trust run by the community. Prior to getting a certificate issued under her name, a user has to get her identity verified. To achieve this, she meets an already well-verified community member, an Assurer. The Assurer verifies her identity documents (this process is called Assurance) and records the fact that he has done so on paper (which he archives) and in the web application. With each Assurance the Assuree gets a number of Assurance Points, based on the documents provided and the experience of the Assurer. When the user has got enough points (which requires at least two Assurances) she can get certificates issued on her name. After further verification (at least three Assurances in total) and a short test on the knowledge about the Assurance process she may become an Assurer herself

[1] Both free as in freedom and as in free of costs

and verify other users. The more Assurances an Assurer has already done, the higher his experience is assumed to be and therefore the more Assurance points he may issue.

When there are disputes or irregularities in the community, such as mistakes made during an Assurance, they are brought into arbitration. In arbitration senior members of the community, called Arbitrators, review the case, taking into account existing CAcert policies and principles, and give a ruling which is binding to all members.

Most of the core functionality of CAcert is handled by a custom software system, which is described in the following sections.

2.2.1. Features of the Existing System

Some of the main features of the existing software are the following:

- Common functionality not tied to a user account
  - Informational web pages
    Display the home page, help pages, etc.
  - Contact support staff
    Send a message to the support team, who can help with account issues and other problems.
- User account management
  - Register new account
  - Change account details
  - Reset password
    Allows to reset the password if it has been forgotten, using lost password questions and an email confirmation for authentication.
  - Organisational accounts
    For companies and other officially registered organisations, with a possibility to have more than one administrator managing the account and including the organisation name in the certificates.
- Web of trust
  - Enter Assurances
    Allows an Assurer to check the information present in the system against the information gathered in the Assurance and allocate Assurance points.
  - Check verification level of a user
    Shows how many Assurance and experience points a user has gathered and how this value is derived.
  - Find Assurers in the Assuree's area
    Helps a new user to find Assurers in her area and contact them in order to meet and get assured.
- Certificate issuing
  - Verify alternative identities (email addresses, domains etc.)
    Allows to manage email addresses and domain names, which will be verified when added to prevent abuse (by sending a verification email to the specified address).

  - Accept certificate requests and generate certificates
    Allows users to submit a certificate request for a previously verified identity and to retrieve the resulting certificate.
  - Remind users of expiring certificates
    Users are reminded to renew their certificates before they expire (one of the most common problems on sites using encrypted connections).
  - Revoke certificates
    If a key might have been exposed or if it was lost, users should revoke the corresponding certificate so that other parties get notified of the problem when they are about to use the certificate.
  - Renew certificates
    Issue a new certificate for the same key pair and with the same details as in the existing certificate but with an updated validity period.

- Administrative functions
  - Access user details
    Search for an account and view its details.
  - Password reset by the support staff
  - Change some user details
    Allows the support staff to change account data like the name and day of birth even after an Assurance already has been issued, which for example might be needed on request of an Arbitrator.
  - View and revoke Assurances
    Support staff may have to revoke an Assurance on request of an Arbitrator or if an Assurer realises he made a mistake just after entering the Assurance.
  - Lock down accounts
    If an account is to be terminated or subject to further investigation in an arbitration case, access to it can be blocked.

2.2.2. Problems of the Existing System

The existing software already implements the basic features but maintaining and reviewing is difficult because the design is of poor quality. The structure loosely resembles some kind of Model View Controller architecture but the components are not well-separated and the model only consists of a plain SQL database with no abstraction on top of it (model classes, object relational mapping or something in that direction). As a consequence of the missing model abstraction, the controllers are huge and even the views have some logic in them apart from presentational concerns. The same pieces of logic (often containing raw SQL statements) appear in many places of the software, so changing behaviour requires quite a few changes.

The software is not well-modularized. The functionality that does not require an user account and the web of trust functionality are separated. The rest, however, is located in one big controller which is over 3000 lines of PHP code, largely without any comments. Even in these large files, there is almost no use of functions or methods and therefore almost no code reuse. As a consequence, the code resembles what commonly is referred to as “spaghetti code”. The software doesn’t use object-oriented programming paradigms and no framework or abstraction library for web application development is used apart from the ones built into PHP. There are only a few libraries used at all and as strings are not encoded in Unicode, proper internationalisation is a major problem.

Another issue for the whole CAcert project is that the CAcert root certificate is not yet included in the major browsers by default. So users who don't know about CAcert or haven't imported the CAcert root certificate into their browser will get a warning message when visiting a website secured with a CAcert issued certificate. In order to solve that problem CAcert has to get audited. That means that an independent Auditor reviews all policies and other documents and checks if the procedures actually implemented in CAcert (both in software as well as manual processes) fit the requirements set by the policies. This Audit also contains a code review of the critical parts of the system. This code review, however, would probably take a lot of effort because of the code quality of the existing implementation, as already described.

All in all the software is in such a bad state that a complete redesign seems unavoidable in order to review and maintain it.

A software workshop held by CAcert developers in April 2009 came to a similar conclusion. Some ideas have been exchanged, some requirements and best practises named and an idea for a design approach has been proposed. This idea however was so rough that it didn't provide enough guidance for the implementation phase that should have followed. There were some efforts to get the implementation going but they died soon after they had started.

2.3. Method for Evaluation

For the evaluation of the designs, the attack tree analysis is used (according to [VM01], based on the fault tree analysis described by [Lev95]). It tries to cover all known attacks possible on the system in a systematic fashion: by grouping them into a tree structure and then assign a probability to them being successful. In this tree structure the root represents the main goal the attacker is trying to achieve. Each child node defines a more detailed possibility to achieve the goal of its parent, eventually resulting in leaf nodes representing concrete attacks. Usually child nodes are to be read as options meaning that attaining any goal of a child node results in the parent goal being achieved (logical OR), but there can also be nodes where all goals of the child nodes have to be attained in order to achieve the parent goal – this is marked by the phrase AND on every but the last child node. When the tree has been completed, a risk value is assigned beginning from the leaves. The risk values of each child node can be cumulated into a risk value for the parent node. For OR nodes this is most likely the risk value of the child with the highest risk while for AND nodes it is the lowest. However if there are many OR child nodes it may be more realistic to assign a higher risk value to the parent to account for the many possibilities to achieve the goal where only one successful attack is needed from an attacker's perspective and vice versa for many AND child nodes.

According to Viega and McGraw the attack tree analysis has the drawback that it is subjective to a certain degree [VM01]. Firstly, the completeness of the attack tree depends on the knowledge about possible attacks, yet unknown attacks are not accounted for, and secondly assigning the risk values is mainly an educated guess where in some cases evidence may be gathered to support it. So in conclusion the accuracy of the attack tree analysis is better when carried out by a team of experts.

The problem is, that measuring the security of a system, especially during the development or even design phase, is difficult. Even more so because objective metrics that are widely-known, accurate, and proven, do not exist in that area. So the attack tree analysis was chosen for its practicality.

2.4. Related Work

Mouratidis et al. [MGM03] try to approach security with a method derived from Tropos covering the whole requirement analysis, software design, and development process. They model the system as a set of related actors, secure entities (goals, tasks and resources), secure dependencies and security constraints. Although their main contribution is to introduce a process that uses the same concepts and notations for the entire development process, the models used in each development stage are very different and can not be easily obtained just by transforming and enriching the diagram from the previous stage. Therefore, the benefit of using the same notation in each development stage can be seen as less significant and choosing a specialised representation for each stage might have advantages. Furthermore most of their proposed process is represented in diagrams, which can get quite complicated, and they have not evaluated their approach.

Starting with Yoder and Barcalow [YB98], various security patterns have been proposed to equip software developers with reusable building blocks to solve security problems [SR01, FP01, SFBH$^+$06]. But they mostly address a single specific problem and thus can be seen as building blocks, not so much as an approach for designing a whole system.

Fielding [Fie00] provides an overview over various architectural styles often used in network-based software, with regard to performance, scalability, simplicity, modifiability, visibility, portability and reliability, but he did not explicitly cover security.

Viega and McGraw [VM01] introduce some basic knowledge about developing secure software and risk assessment, including the attack tree analysis used in this work. They stress the importance of doing a security analysis in the design phase to find flaws on the architectural level, which are hard to fix later on without introducing new flaws, but they do not explicitly cover architectural styles.

In general, there are many approaches that propose solutions to recurring problems, as in the pattern community, or even a whole design process, but a systematic or even formal evaluation is often missing. Also the expectation, that there will be errors in the implementation, and one of the goals of a good design with security in mind is to confine the impact of these errors, is often not reflected in the description of the patterns and processes.

3. Approach

This work is structured in three phases, as depicted in Figure 3.1: First, there is a detailed requirement analysis, which apart from the functional requirements also contains non-functional requirements (including security) and a risk assessment. The second phase is an iterative evaluation of various architectural styles, followed by the third phase in which the styles from the previous phase are compared, to find out whether the choice of the architectural style bears any significance and which one is the most suitable for the example system. A detailed description of each phase is given in the following sections.

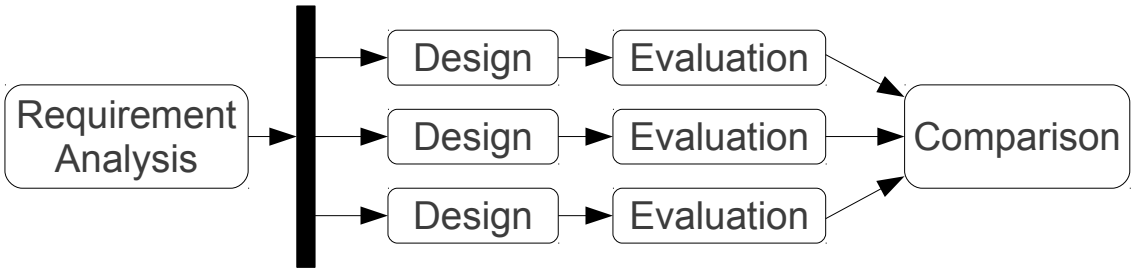

Figure 3.1.: Schematic Illustration of the Approach

3.1. Requirement Analysis

There was no existing requirements document describing all functional and non-functional requirements for the system. As a thorough understanding of the requirements provides the base for the design, a requirement analysis was conducted.

The requirements were derived from various sources: Most of the functionality of the existing system should be preserved except where good arguments are provided why the functionality is disadvantageous, unnecessary, or should be implemented outside the scope of the system. Additionally, there exist quite a few policies from CAcert or affecting CAcert that also have implications for the software. These should be respected, although policies are not immutable and if there is a good reason raised in the design process, they may be changed (at least the CAcert-driven ones). Another source of requirements were various ideas about new features that could not easily be implemented in the existing system. Some of those ideas were documented in the previous design attempt mentioned in Section 2.2.1, others were mentioned in bug reports or posts on mailing lists discussing CAcert. Finally there are also many domain experts in the community (including the author of this thesis) who were asked for additional requirements not yet covered by the other sources.

In addition to the current functional and non-functional requirements for the system, the requirement analysis also tried to cover future requirements, i.e. requirements that are currently not present but where it is probable that a change of requirements may be needed in the future.

Adding directly to the requirement analysis a risk assessment was performed (roughly following [SFBH+06]). First, the main critical assets were identified. Then, the needed

kind of protection for each asset (confidentiality, integrity etc.) as well as the estimated damage if this protection fails were identified. The result is a document which can be found in Appendix A that also codifies some security requirements or can be used to extract them.

The resulting requirements document is rather detailed to facilitate the design process that followed. The highly detailed specification of the requirements also helped to discover discrepancies in the requirements themselves and probably avoided problems in subsequent development stages.

3.2. Evaluation of Architectural Styles

Following the requirement analysis, several architectural styles were evaluated in order to be able to compare them. The process for evaluation was as follows.

First, a style was selected for evaluation and a rough assessment just on the known properties of the style was done, for example in regard to whether it is suitable to fit the functional and non-functional requirements or whether it has unacceptable implications. Afterwards a rough design following the style (if it wasn't already classified as unsuitable) was produced, taking into account the functional and non-functional requirements and incorporating safeguards against assumed security risks. The resulting designs can be found in Chapter 4.

Then the designs were evaluated based on the attack tree method and "common sense". The results of the evaluation and a more thorough description of the evaluation process is provided in Chapter 5.

3.3. Comparison of the Architectural Styles

The evaluations from the previous phase were studied to find differences in the risk values. Where differences were found it was tried to retrace where they stemmed from, for example if there were any structural differences, implied by properties of the architectural style, that caused the differences in the risk evaluation. This required additional work on the existing designs of the compared styles if the rough nature of the designs didn't allow for a detailed enough analysis. Based on this comparison it was determined if the choice of the architectural style had a significant impact on the attack risk of the resulting design.

In addition to this comparison, a decision was made which design was the most suitable for the example system. This decision did not only take security requirements but all functional and non-functional requirements into account.

4. Architectural Design Alternatives

This chapter describes the architectural styles selected for evaluation and the design that was created for each style. For each style that was considered a description of the style is given and it is explained why the architectural styles layered architecture and service-oriented architecture were chosen for evaluation and why the pipes and filters architecture was not pursued further. Also a detailed description of the resulting design according to the two selected architectural styles is given.

4.1. Layered Architecture

The first architectural style selected to be evaluated was the layered architecture style. A layered architecture is characterised by the separation of components into different levels of abstraction, called layers. The upper layers use the lower layers to provide functionality, but never the other way around. When using a strict layered architecture a component may only use a component in the layer directly beneath it while in a relaxed layered architecture a component may use components in any lower layer (i.e. skip a layer). Layering may be enforced by using runtime protection mechanisms such as the CPU privilege levels (utilised in protection of operating systems) or separating the layers in tiers, which means that the layers are grouped and each group is executed in a separate process or even on different machines, which interact in a client-server style manner, with the lower layer being the server for the upper layer. A more detailed description of the layered architecture pattern can be found in [BMR+96].

4.1.1. Choice

The layered architecture pattern was selected because it is often used in contexts where security plays a major role. For example operating systems are often structured in layers, separating the user space from the kernel space and possibly have further layers within the operating system kernel. Also quite a few typical business applications are designed as a multi-tier architecture.

4.1.2. Design

The system is divided into five layers using strict layering (i.e. no layer may be skipped). Strict layering was used in order to confine the effects of an attack against upper layers. So an exploit of one of the upper layers can not be used to directly access and attack the lowest layer. The resulting architecture design can be found in Figure 4.1. The following sections describe the different layers in more detail.

4.1.2.1. Front-End Layer

The topmost layer is the front-end layer and consists of all components directly facing the external network. These components are further divided vertically into the web interface, the login server, and remote APIs for third parties. The authentication and credential management components were separated from the rest to isolate them, assuming that the code base may be rather small and better reviewed, resulting in a smaller attack surface

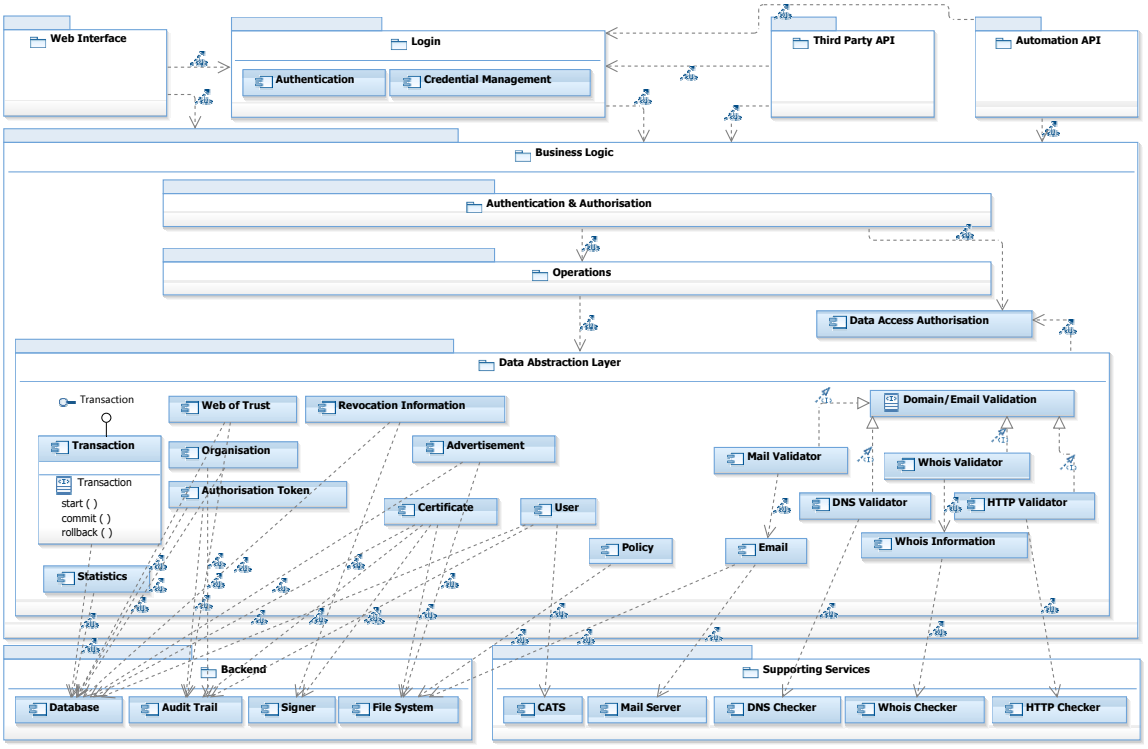

Figure 4.1.: Layered architecture design – larger version in Appendix B

on the login credentials. Other components will not directly work on the credentials and only get an authentication token from the client, which they can use to verify that the client has been successfully authenticated by the authentication component by querying the authentication and authorisation layer. Other front-end components still have to check the re-authentication credentials though. The front-end components realise all functionality they can not handle by themselves by making a request to the lower layer, forwarding the authentication token for proof of authentication.

The front-end components do not keep any state so they can be scaled up by deploying components multiple times, without needing to migrate state between machines, and it gets more difficult to expose restricted information from another user by attacking the front-end components. If there is interaction context that can not be held by the client (by using cookies or request parameters), which for example might be the case with cross-site request forgery (CSRF) protection mechanisms, it should be forwarded for to lower layers for storage.

4.1.2.2. Authentication and Authorisation Layer

The second layer is the authentication and authorisation layer. The authentication state of a client is forwarded to this layer by the authentication component in the front-end layer on successful authentication of the client, including all information needed to validate the authentication token provided by the client on successive requests. For each request coming from the front-end layer, the validity of the authentication token and whether the authenticated user is authorised to execute operation requested is checked. If the authorisation is granted, the request is forwarded to the operations layer, including information about the authentication state.

The authentication and authorisation layer should only keep soft state, meaning state that is recoverable and that is discarded after some time-out or on demand. This allows some possibly expensive to reproduce state to be cached and still makes it possible to deploy

the layer on multiple machines. When the layer is deployed multiple times, a mechanism should be used that allows to invalidate certain state in all deployed copies, to allow to safely erase authentication state on logout.

4.1.2.3. Operations Layer

The operations layer offers high level operations possibly involving multiple data objects. It allows to orchestrate multiple data objects, without having to encode complicated choreographies, not really belonging to one object or another, into the data objects resulting in bloated data objects and too tight coupling. As a result of the strict layering and authorisation checking on operation level, the operation layer also has to provide view operations forwarding properties of data objects up to higher layers. However these view operations may bundle data objects often requested at once in one single operation, reducing the total number of requests needed by upper layers and thereby improving performance. Authentication state information also has to be forwarded to the data access authorisation component in the data abstraction layer.

The operations layer should not keep any state and therefore may be replicated, but care should be taken that the transaction component of the data abstraction component is used where required.

4.1.2.4. Data Abstraction Layer

The data abstraction layer mainly consists of data abstraction objects, providing an object oriented encapsulation to the underlying data assets, like records in a database, certificates and other information. These data abstraction objects also provide methods manipulating the objects, simple operations concerning mainly one object and some constraint checking. This constraint checking at data object level involves for example verifying that there are no valid certificates left when removing a domain, but also checking authorisation on data object level. This authorisation checking is done through the data access authorisation component (which is part of the data abstraction layer). The component uses the information about the authentication state that is forwarded from the authentication and authorisation layer, through the operations layer, to the data abstraction layer, to check whether the particular user may access the object in the requested way.

As duplicated data objects may result in data inconsistencies, even when using transactions (due to higher risk of mistakes in implementation), the data abstraction objects should not be replicated. It would be possible however to partition the data objects to spread the load.

4.1.2.5. Back-End/Supporting Services Layer

The back-end consists of components that actually store the data represented by the data abstraction objects. That includes data bases, as well as the file system, the audit trail server, and the signing server (for issuing and revoking certificates). The supporting services are components that connect to external services on behalf of the data abstration objects. The data abstraction objects will never directly connect to an external service (for example to verify a domain name), they will always use a component operated by CAcert to act as an intermediary. These may be off-the-shelf components, such as in the case of a mail server which delivers the locally submitted emails to foreign mail servers, or custom components written for this purpose.

4.2. Pipes and Filters Architecture

The pipes and filters architectural style was evaluated for suitability to fit the requirements. Each request would be interpreted as an object, routed through a system of pipes and filters, eventually being transformed into a response. The problem was that, in order to process the request, quite some filters would add information to the each request, which would then be consumed by later filters. This approach results in many stages of filters and quite some pipes, meaning that data would have to be passed around and possibly be copied quite a lot which makes it inefficient. Also it would result in an unintuitive architecture, which might cause problems in development. Therefore the architecture was not pursued any further.

4.3. Service-Oriented Architecture

The last architectural style to be evaluated was the service-oriented architecture. A service oriented architecture is characterised by autonomous components, called services, which provide an interface via an implementation technology independent protocol. These services are combined to provide the functionality of the complete application. By using a protocol that is independent from the implementation, the services can be loosely coupled and implemented in any technology.

4.3.1. Choice

The service-oriented architecture pattern was chosen because the services can be fully independent, allowing for vertical partitioning of the system, which confines the effects of an attack, and because it is a pattern commonly found in modern business applications.

4.3.2. Design

In the design depicted in Figure 4.2 the system consists of front-end components, the business logic, an authentication and authorisation service, and supporting services. The general idea is that each service is self-contained, meaning that for example it does not rely on another service for data storage, and cooperates with other services in a choreography. There is no managing component coordinating the cooperation between services. This is different to the layered architecture design, where higher levels mostly rely on lower layers to provide essential operations, such as persisting data, and operations involving multiple data objects are coordinated by the operations layer.

4.3.2.1. Front-End

The front-end components are very similar to the components in the front-end layer in the layered architecture design (Section 4.1.2.1), except that interaction context that can't be held by the client is stored in the front-end component itself. This makes it necessary to synchronize the state if the front-end component is deployed on multiple machines.

4.3.2.2. Authentication and Authorisation

The authentication and authorisation service is responsible for authentication of clients, managing credentials, and is queried by other services for the authorisation status for each invoked operation. Like in the layers design, authentication tokens are used to keep track of the authentication status. Therefore, they have to be supplied for each invoked operation and are in turn provided to the authentication and authorisation service.

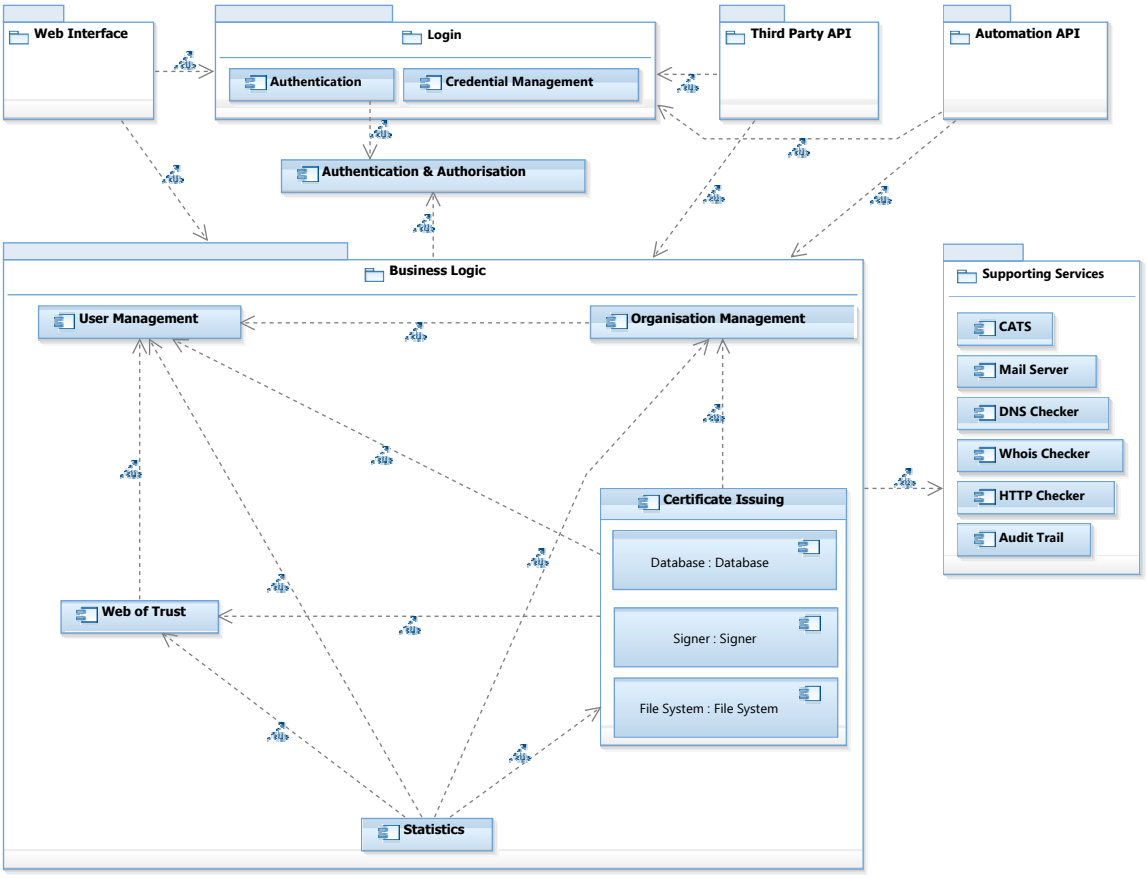

Figure 4.2.: Service-oriented architecture design – larger version in Appendix D

4.3.2.3. Business Logic

The business logic consists of some services providing the core functionality. The services already contain what has been in the back-end in the layers design and connect to supporting services as found in the layers design where needed (see Section 4.1.2.5). As already indicated, the collaboration of the services is not managed by another component, but each service is responsible to collaborate with other services of its own accord. For example if the user management service gets a request to remove a domain from the account, it should also invoke the certificate issuing service to revoke all valid certificates issued for that domain. As the services also are responsible for persisting their data, they can't be deployed on multiple machines without some preparation.

5. Evaluation

In this chapter the designs presented in Chapter 4 are evaluated and compared with each other. First, in Section 5.1 the variation of the attack tree method used in this work is described. Afterwards, in Sections 5.2 and 5.3, the evaluations of both designs are presented and then compared in Section 5.4. The evaluation shows that there are structural differences in the risk evaluation but no significant differences in the result, even when ignoring a common dominating specification-inherent risk.

5.1. Variation of the Evaluation Method

The attack tree analysis presented in Section 2.3 is annotated with risk values. Viega and McGraw assign a single risk value to each node that is derived from the estimated time (effort), cost and risk to the attacker [VM01]. This assessment is slightly modified for this work. Firstly the components making up the risk value are not merged but kept separate, to make it clearer how the risk is derived and make the propagation up in the tree more precise. Secondly time and cost are tightly related in that attacks can usually accelerated if more money is invested and missing capital can be compensated to some point by investing more time, therefore they are combined into a single value called effort. Thirdly a component called gain is added, which represents the motivation for the attacker to carry out the annotated attack – not only in the context of the current goal but in general. This is based on the assumption that an exploit that is useful in many contexts attracts more attackers (e.g. it is more valuable to find an exploit in a standard database component or even an encryption protocol than in a custom component used in a single system) and an exploit that results in more capabilities in itself is more valuable for the attacker than one that doesn't and therefore is more likely to be available in any context. The scale of the components ranges from 1 (meaning low effort, risk or gain) to 9 (meaning high effort, risk or gain).

5.2. Layered Architecture

This sections presents the evaluation of the layered architecture design from Section 4.1.2.

5.2.1. Security Evaluation

An evaluation of the attack risk was conducted according to the attack tree method presented in Section 5.1. The completed attack tree resulting from that evaluation can be found in Appendix C.

The result shows, that the risk to the web of trust, the user data, and the confidentiality of the issued certificates, is dominated by the ability of users to keep their login credentials secret (Section C.2.2). This means that the main risk identified lies in the specification/environment of the system and not in the system itself. One solution would be to disallow the use of passwords as means of authentication, but this would violate the requirements and make the system unusable for some users. As some assets can only be accessed by privileged users (e.g. the user data may only be modified by support engineers if the account already got assured), one could argue that these users take extra care with

their login credentials and the effort for the attacker is higher in these cases, this however doesn't solve the general problem. A mitigation worth exploring, might be requiring two-factor authentication whenever passwords are used (e.g. by sending an additional token to the primary email address). But this might have severe consequences if the additional authentication factor fails (for example if the control over the primary email address is lost), because then, even the legitimate user might not be able to log into the account to change the authentication settings.

As the domination by a specification-inherent risk makes comparison between the risk evaluations of different styles less meaningful (they would all show the same result where the domination takes place), it was investigated how the attack tree would change if the dominating risk was assumed to be mitigated in some way. The result is, that now the hijacking of the authentication state (Section C.1.1 Item 4.3.3.4.) would become the dominant risk, leading to an effort value of 4 and a value for risk to the attacker of 2, for the web of trust, the integrity of login credentials, and other analogous cases. The risk to the confidentiality of the login credentials is dominated by the risk of intercepting the credentials on the login server (Section C.2.2 Item 5.), which results in a value for effort and risk to the attacker of 5 and 1 respectively.

5.2.2. Overall Evaluation

The design was created with scalability in mind, meaning that the upper layers can be replicated onto multiple servers if needed. The components in the back-end may be replicated, depending on their support for this use case. Maintainability and testability should also be good, because of the rather stable interfaces provided by the layers, which confine the effect of code changes [BMR+96, Fie00], and the possibility to emulate lower layers by mock objects. One drawback might be high overhead, caused by the redundant processing of data in each layer [BMR+96, Fie00], which may even be increased if the layers are deployed in tiers, which means additional overhead for serialising and transferring the data between different address spaces or even machines.

5.3. Service-Oriented Architecture

This sections presents the evaluation of the service-oriented architecture design from Section 4.3.2.

5.3.1. Security Evaluation

Like for the layers design, an evaluation of the attack risk, according to the attack tree method presented in Section 5.1, was conducted. The resulting attack tree can be found in Appendix E.

Unsurprisingly the result shows that, like in the evaluation of the layers design (Section 5.2.1), the risk to the web of trust, the user data, and the confidentiality of the issued certificates, is dominated by the ability of users to protect their login credentials (Section E.2.2). If one assumes, for the sake of being able to compare the designs, that the dominating risk of users not being able to protect their credentials would somehow be mitigated, an attacker exploiting the operation logic of a service becomes the dominating risk (Section E.1.1 item 3.1.). Accordingly, the effort the attacker has to make increases to 4 for the web of trust, the integrity of login credentials, and other analogous cases. For the confidentiality of the login credentials the risk of intercepting the credentials on the login server becomes the dominating risk (Section E.2.2 item 5.), leading to an effort value of 5.

The finding that the weakness of the design is that an error in a part of the business logic service may compromise the whole service, shows that the expectation when choosing the architectural style, that the vertical partitioning might help to confine the effects of an attack, was not entirely met. One cause might be, that the scope of a service is too broad, so that the attack surface (which is all operations offered by that service) is too big and can't compensate for the missing horizontal partitioning. This might be mitigated by narrowing the scope of each service, so that only few operations are dependent on the security of each other, but that is only possible to a certain degree. One limiting factor here, is the data shared by the operations. If two operations working on the same data are split into separate services, then still both services have to access the data somehow and an attack on one of these services gives access to that data shared by both. So this approach only works if the data sets the services operate on still differ in some way, and it has the drawback that invocations that were previously internal to the service become part of the external interface, which especially with service-oriented architectures comes with some overhead, as each request has to be translated into the protocol and then transferred, and also might degrade the cohesion of the component.

One concern resulting from the decoupled structure of the service-oriented architecture, is that an attacker might exploit the connection between the components instead of attacking the components themselves. However this risk is negligible compared to the risks laid out above.

5.3.2. Overall Evaluation

Due to the principle of every service keeping its own data, it is assumed that deploying a service on multiple machines is generally not possible and requires the service to be built with some sort of synchronisation mechanism. So scalability is not implied by the design and requires some additional work. Maintainability should be good, because the services have well-defined interfaces, which confines coding changes, and the use of an implementation technology independent protocol allows for exchanging a service with another, implemented using a completely different technology. Testability is to a large part dependent on the internal design of a service. Though tests on the service interface level are possible, they are on a rather rough granularity. The communication overhead of the architecture is expected to be quite significant, because each request has to be transformed into the standard protocol and then transferred to the other service, and each operation involves a couple of services (a front-end component, the authentication and authorisation service and at least one business logic service). Due to the decision not to have a component coordinating the cooperation between services (choreography instead of orchestration), there might be situations where there are redundant operations (e.g. when deleting a user account, the user management service might request the certificate issuing service to revoke all valid certificates associated with the account and then call its own procedure to delete all domains associated with the account, resulting in another request to the certificate issuing service requesting to revoke all valid certificates issued to that domain, although they already have been revoked). However it is expected that in comparison the overhead resulting from that issue is rather small.

5.4. Comparison

Even when ignoring the dominating risk of users not being able to protect their login credentials, the differences in the result of the security evaluation of both designs are marginal (there is a difference in risk to the attacker of one point between the two designs) and are within the accuracy of the method used for evaluation. If anything, the only indicator that can be used is the kind of dominating risk. For the service-oriented architecture design the

dominating risk is an attacker being able to exploit the operation logic of a service, which indicates a weakness in the system, while the dominating risk of the layers design lies in the attacker hijacking the authentication state of a user, which is an attack outside the system. When looking at the dominating attack on the service-oriented design, one can see that in the layers design the same attack is limited due to the authorisation checking on the data abstraction layer. So even if an attacker is able to exploit an operation to manipulate data objects in unexpected ways, he still needs to pass the data authorisation checking. On the other hand, when looking at the dominating attack on the layers design, the same attack is also possible in the service-oriented design. So there might be a slight advantage of the layers design from the security point of view, but it really is small.

The maintainability should be about the same for both designs. One aspect where the service-oriented architecture is usually better suited, is if one wants to integrate components implemented using different technologies, as is usually the case with legacy components or components developed in different departments of an organisation (or even not implemented within the organisation at all). However, this use case is probably not needed within CAcert, because there is no part of the legacy application suited for reuse, and CAcert as an organisation is not so big, that it would have different departments doing implementation work independent from each other. Testability is probably a little bit better with the layers design, because of the stable interfaces in the vertical direction, allowing to confine the object under testing. If the layers design is deployed in tiers, the communication overhead is probably higher than in the service-oriented design, because each request path has to pass more boundaries where data is serialised and transferred (at least one time from the topmost layer to the lowest and back again). However the layers design is likely more scalable, as the upper layers can be deployed redundantly while keeping the synchronisation overhead low.

All in all the arguments for the layer design seem stronger in this particular environment and it is therefore recommended for implementation.

6. Conclusion

In the evaluated scenario a significant influence of the architectural style on the security of the resulting system could not be substantiated. There were some small differences in the resulting risk of certain attacks, but they were too small to distinguish them from possible inaccuracies resulting from the method of evaluation. This does, however, not prove that such an influence does not exist, just that in the evaluated environment there was not enough evidence to prove the existence of such an influence. It was tried to establish rather strict conditions to make a potential positive result stronger and more realistic. The architectural styles examined in detail in this work were both chosen for their suitability in security contexts known a priori, if other styles would have been evaluated the evaluation might as well have resulted in a different outcome.

What has been presented in this work, is that a security evaluation can be integrated into a very early stage of the design process. Those security evaluations can show possible attack risks that should be addressed in further design iterations. Integrating security reviews this early in the design process might help to fight the problems that come with the "security as an afterthought" approach.

Because mutliple designs were evaluated in this work, solutions that address certain problems in the earlier designs have been reused in the following design processes. This makes it hard to tell, whether there was an advantage in efficiency in the design process for one architectural style or another. Future work could examine this question, for example by designing an experiment measuring the time taken for the design process and security of the resulting system when using different architectural styles.

Another interesting question for future investigation would be, whether an influence of the architectural style can be shown for more relaxed environments, which still seems likely, or if it can be universally shown to be absent.

Also a different selection of the designs to be evaluated could be promising. For example, evaluating existing systems with the same specification and different architectures for their expected risk for attack, using the method presented in this work, and comparing the result to their real historic security record. Candidates for such an evaluation might be standard components that are regularly scrutinised for their security properties, such as SQL data bases or web servers.

Bibliography

[And01] R. J. Anderson, *Security Engineering – A Guide to Building Dependable Distributed Systems*. Wiley, 2001.

[BMR+96] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, and M. Stal, *Pattern-Oriented Software Architecture: A System of Patterns*. John Wiley and Sons, New York, 1996, vol. 1.

[CCCN99] Committee on Information Systems Trustworthiness, Computer Science and Telecommunications Board, Commission on Physical Sciences, Mathematics and Applications, and National Research Council, *Trust in cyberspace*, F. B. Schneider, Ed. National Academy Press, 1999.

[CM07] S. Christey and R. A. Martin, "Vulnerability type distributions in CVE," *Mitre report, May*, 2007. [Online]. Available: http://cwe.mitre.org/documents/vuln-trends/vuln-trends.pdf

[Com10] Committee on National Security Systems, "National Information Assurance (IA) Glossary," Apr. 2010. [Online]. Available: http://www.cnss.gov/Assets/pdf/cnssi_4009.pdf

[DR08] T. Dierks and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2," RFC 5246 (Proposed Standard), Internet Engineering Task Force, Aug. 2008, updated by RFCs 5746, 5878, 6176. [Online]. Available: http://www.ietf.org/rfc/rfc5246.txt

[Fie00] R. T. Fielding, "Architectural Styles and the Design of Network-based Software Architectures," Ph.D. dissertation, University of California, 2000. [Online]. Available: http://jpkc.fudan.edu.cn/picture/article/216/35/4b/22598d594e3d93239700ce79bce1/7ed3ec2a-03c2-49cb-8bf8-5a90ea42f523.pdf

[FP01] E. B. Fernandez and R. Pan, "A pattern language for security models," in *Conference on Pattern Languages of Programming (PLoP 2001)*, 2001. [Online]. Available: http://hillside.net/plop/plop2001/accepted_submissions/PLoP2001/ebfernandezandrpan0/PLoP2001_ebfernandezandrpan0_1.pdf

[ITU11] ITU, "X.509," Feb. 2011. [Online]. Available: http://www.itu.int/rec/T-REC-X.509

[Lev95] N. G. Leveson, *Safeware – System Safety and Computers*. Addison-Wesley, 1995.

[MGM03] H. Mouratidis, P. Giorgini, and G. Manson, "Integrating Security and Systems Engineering: Towards the Modelling of Secure Information Systems," in *Advanced Information Systems Engineering*, ser. Lecture Notes in Computer Science, J. Eder and M. Missikoff, Eds. Springer Berlin / Heidelberg, 2003, vol. 2681, pp. 1031–1031. [Online]. Available: http://dx.doi.org/10.1007/3-540-45017-3_7

[PW92] D. E. Perry and A. L. Wolf, "Foundations for the Study of Software Architecture," *ACM SIGSOFT Software Engineering Notes*, vol. 17, no. 4, pp. 40–52, 1992. [Online]. Available: http://dx.doi.org/10.1145/141874.141884

[SFBH+06] M. Schumacher, E. Fernandez-Buglioni, D. Hybertson, F. Buschmann, and P. Sommerlad, *Security Patterns: Integrating Security and Systems Engineering*. Wiley, 2006.

[SR01] M. Schumacher and U. Roedig, "Security Engineering with Patterns," in *The Proceedings of the 8th Conference on Pattern Languages of Programs (PLoP)*, 2001. [Online]. Available: http://hillside.net/plop/plop2001/accepted_submissions/PLoP2001/mschumacher0/PLoP2001_mschumacher0_1.pdf

[Sta11] W. Stallings, *Cryptography and Network Security – Principles and Practice*, 5th ed. Prentice Hall, 2011.

[VM01] J. Viega and G. McGraw, *Building Secure Software: How to Avoid Security Problems the Right Way*. Addison-Wesley, 2001.

[YB98] J. Yoder and J. Barcalow, "Architectural Patterns for Enabling Application Security," in *Conference on Pattern Languages of Programming (PLoP 1997)*, 1998. [Online]. Available: http://hillside.net/plop/plop97/Proceedings/yoder.pdf

Glossary

architectural style (sometimes also called "architectural pattern") A common design that is applied to the whole architecture of a system. It often specifies how components of the system interact with each other.
Examples of architectural styles: Pipes and Filters, Client/Server, Layered Architecture, Service-Oriented Architecture. iii, 1, 7, 9–11, 14, 19, 21

Assurance The process by which the Assurer verifies the identity of an Assuree. Sometimes it also refers to the data record created in the web application for an Assurance. 3–5, 30–32, 35–37, 39–42

Assurance points For each verification of the identity of an Assuree the Assurer assigns a confidence value, which is expressed in Assurance points. The amount of Assurance points the Assurer may issue depends on his experience. 30, 32, 33, 35, 36, 39, 41, 42

Assuree CAcert community member that wants to get her identity verified. 3, 4, 36, 37

Assurer CAcert community member that is already well-verified. He will check the documents of Assurees to verify their identity. 3, 4, 27, 31, 32, 35–38, 40–42, 44

CA certification authority. iii, 1, 3, 43

CERT (computer emergency response team) often refers to the CERT Coordination Center located at the Carnegie Mellon University, which issues advisories to system administrators and developers about common existing or potential threats to computer security. 1

CSRF (cross-site request forgery) is an attack where malicious data is placed on a third party website that causes a user agent displaying that third party website to make a cross-site request to the target website possibly without the user knowing about the request being made. This forged request includes authentication state such as cookies if the user has been previously logged in to the website and may lead to side effects which seem to happen on behalf of the user such as a transaction being made, an item being purchased, a message being sent etc. 12

CVE (common vulnerabilities and exposures) is an index of publicly known information-security vulenarabilities and exposures. 1

TLS (transport layer security) specified in [DR08] is a successor to the SSL (Secure Sockets Layer) protocol. It provides an authenticated, confidential connection oriented transport service for application layer protocols. It is best known for its use with HTTP to securely deliver web sites. 3

X.509 A standard format for certificates [ITU11]. It contains identity information about the subject identified and its public key and can be extended in various ways. Many services use X.509 certificates to establish the identity of the other party involved in communication, most noteworthy TLS and S/MIME (a format for securing email communication). 3, 28, 33–35, 38

A. Requirements

A.1. Functional Requirements

1. Functionality not tied to an user account

   1.1. Static web pages

      1.1.1. Root certificates — Existing Feature

         1.1.1.1. HTTP download in various formats (PEM, DER, PGP)

         1.1.1.2. Import functionality
         Some browsers need extra functionality for easy import of the root certificate (especially Internet Explorer)

         1.1.1.3. Download of the current CRLs
         The CRLs might be partitioned in the future to make them smaller in size, so multiple CRLs for the same root have to be supported

      1.1.2. Policies — Existing Feature
      The system should provide a repository of normative policies. This might be moved out of the critical system area in the future. This repository should include the state of the particular policy (work in progress, draft, policy etc.)

      1.1.3. Statistics — Existing Feature
      The system should provide some statistics about itself (issued certificates, Assurers etc.). The statistics may be regularly pre-generated

   1.2. Contact Support Team — Existing Feature

      1.2.1. List of various places to ask for help
      Link to such a list in a non-critical system (e.g. the wiki, CMS)

      1.2.2. Contact form
      This might be moved to a non-critical system (e.g. CMS)

         1.2.2.1. Select where to send the message
         The user should be able to choose one of various places the message is sent to: Support Engineers for sensitive information and things that can only be done by privileged users, support mailing list for general questions. Other places might be added in the future

         1.2.2.2. Spam protection — New Feature

         1.2.2.3. Authorisation token if logged in — New Feature
         If the user is logged in and sends a message to the Support Engineers include an authorisation token as described in requirement 10.2. and a link to the user account in the message

Altered Feature

1.3. Send Emails
Emails sent by the system should include some easy to configure text in every message (such as donation campaigns) and should be signed with a X.509 certificate that is also used by the support team (this distributes the certificate to many users email clients so they can contact us encrypted).

1.4. Internationalisation

Existing Feature

1.4.1. Strings and other resources translatable
Exceptions possible for interfaces only used by a few persons, such as the Support Engineer interface, where it can be assumed that all persons are capable of understanding English. It has to be able to work with our localisation system (that is Pootle which can handle Gettext, XLIFF, Qt TS, TBX, TMX, Java Properties, Mac OSX strings, PHP arrays and some other formats)

New Feature

1.4.2. Unicode
Wherever possible data should be transferred and stored in Unicode aware encoding

Altered Feature

1.4.3. Cultural differences should be respected

1.5. Accessibility

Existing Feature

1.5.1. Unobtrusive JavaScript
All features should be accessible without JavaScript or similar techniques enabled. Only convenience features may only work with those technologies turned on.

New Feature

1.5.2. Accessibility for People with disabilities
All features should be accessible for people with disabilities (e.g. blind people). Where this is not possible alternatives should be available (e.g. sound captchas in addition to visual ones)

1.6. Revocation Status Checking
Third parties need the possibility to obtain the revocation status for certificates issued by CAcert. There are two major techniques for this purpose:

Existing Feature

1.6.1. CRLs
CRLs are files that contain a signed list of the serial numbers of revoked certificates. They may be cached by third parties for various reasons within their validity period (for availability or scalability reasons they may also be cached within CAcert but the cache life time should be drastically lower than the validity period of the CRLs). The authoritative CRL should be updated very often (e.g. once every 15 minutes) in order to keep the time it takes to propagate the revocation state of a certificate low. In order to keep the size of CRLs low, the CRLs may be partitioned (e.g. by month the certificate in question was issued on) but there should be a "master CRL" available which contains the revocation status for all or a significant amount (e.g. all that haven't expired yet) of issued certificates. This is needed for some configurations (e.g. Apache web servers accepting client certificates for authentication).

Altered Feature

1.6.2. OCSP
OCSP is an interactive protocol that allows a client to query the revocation status from a responder by providing a serial number of the certificate and getting back a signed response valid for a time span specified in that response. If the system does not implement a distributable system of OCSP

responders on its own, it has to provide some interface so a separate responder can get the information it needs (at least the revocation status if the serial is provided and some way to successively get a list of all valid serial numbers to be fetched during a very long time span)

1.7. Advertisements — New Feature
It should be possible to include static advertisements hosted at the system in predefined places (plain images or text only, no scripts or interactive objects). Also some statistics interesting for advertisers should be collected (views, clicks etc.; should be modifiable)

1.8. User registration — Existing Feature

1.8.1. Collect user data

1.8.1.1. One or multiple names — Altered Feature
Some guidance should be given how these are to be entered.

1.8.1.2. Date of birth

1.8.1.3. Acceptance of some agreements
At the moment only the CCA (CAcert Community Agreement), there may be more obligatory agreements in the future

1.8.1.4. Email address

1.8.1.5. Credentials for one or more login methods (requirement 2.1.1.) — Altered Feature

1.8.1.6. Credentials for one or more account recovery methods (requirement 2.2.) — Altered Feature
If that is required by the recovery method

1.8.1.7. Opt-in for announcements/newsletters

1.8.2. Verify collected data

1.8.2.1. Verify email address using one or more verification methods (requirement 2.3.4.)

1.8.2.2. Test whether chosen login method works
e.g. for password enter it two times, for challenge/response schemes try one round

2. Normal user accounts

2.1. Login

2.1.1. Several login methods — Altered Feature
At least password and client certificate authentication should be offered. Adding of further methods such as one time passwords, hardware tokens etc. should be possible. Also selecting multi-factor authentication should be possible. Which combinations of login methods and the concrete security parameters for each method are acceptable is subject to varying needs and need to be easily adjustable.

2.1.2. Continuous authentication — New Feature
When authentication methods where this possible without user interaction are used (e.g. client certificate login) the session should be tied to those methods and require reauthentication on every request to hinder session stealing. Also it should be possible to tie the session to a fixed IP (disabling should be possible to allow for the use of anonymisers such as TOR, but maybe only if authentication methods that allow for continuous authentication are used)

Altered Feature
2.2. Account recovery
Several methods for account recovery should be offered. It should be possible to deactivate some methods on the own account. Some of them need to be combined to recover the account. Which ones may be/have to be combined should be easily changeable

New Feature
2.2.1. Password Recovery with Assurance
An already in use mechanism currently manually handled by the support engineers. It uses the Assurance process to provide an out of band channel with authentication (by passports and ID cards). The procedure is described in
`http://wiki.cacert.org/Support/PasswordRecoverywithAssurance`

Existing Feature
2.2.2. Control of the primary email address
The email verification methods (requirement 2.3.4.) should be used to verify control of the main email address

Existing Feature
2.2.3. Reset questions
A set of questions about stable facts about the user that are nevertheless not known to others and their answers are configured. On recovery the user has to answer some subset of them correct (e.g. 3 out of 4). Some generally good questions should be provided (e.g. name of first pet), but the user may add her own.

2.2.4. Other Methods
May be added over time.

2.3. Edit account data

New Feature
2.3.1. Add names
The same restrictions as for registration (requirement 1.8.1.1.) apply

New Feature
2.3.2. Remove names
Names may be marked as deleted at any time, but that also causes certificates in that name to be revoked and the name will only really be deleted when the last certificate expires (even if it already has been revoked). At least one name always has to stay in the account (not necessarily the same name)

Existing Feature
2.3.3. Edit date of birth
Is only allowed if no name on the account has any Assurance points

Altered Feature
2.3.4. Add email address
Once an email address is added to the account, the ownership has to be verified by passing at least two checks described below (required by CPS). The system should allow for the checks to be repeated if later decided that this is needed (e.g. once a year, on every issuing of a certificate if last check too old). If the email address is already assigned to a different account there should be a possibility to transfer it from the old to the new account even if the user has no access to the old account (e.g. do the email checks, then send an email to the old accounts primary email address asking to immediately confirm or refuse the transfer within a fixed time span)

Existing Feature
2.3.4.1. Email ping
A random string is sent in an email to the address to be verified. This string has to be entered in the system (manually or by clicking a link) and then the user should confirm that he wants to add the email address to the specified account

2.3.4.2. Email is used in Assurance (requirement 3.1.) — New Feature
That means the Assurer has a signed paper form stating that the email address is controlled by the user ⇒ legal instead of technical means

2.3.4.3. Other Methods — New Feature
May be added over time. The system should allow for requiring more methods to verify the email address if policy changes in the future

2.3.5. Remove email addresses — Altered Feature
Certificates containing that address should be revoked, address should only be marked as deleted until the last certificate expires. At least the primary email address has to remain on the account

2.3.6. Select primary email address — Existing Feature
One email address is the "primary" email address where notifications, announcements and newsletters get sent to

2.3.7. Add DNS domain name — Altered Feature
Once a domain name is added to the account, the ownership has to be verified by passing at least two checks described below (required by CPS). The system should allow for the checks to be repeated if later decided that this is needed (e.g. once a year, on every issuing of a certificate if last check too old). If the domain is already assigned to a different account there should be a possibility to transfer it from the old to the new account even if the user has no access to the old account (e.g. do the domain checks, then send an email to the old accounts primary email address asking to immediately confirm or refuse the transfer within a fixed time span)

2.3.7.1. Email ping — Existing Feature
A random string is sent in an email to a privileged email address on that domain. Privileged addresses are those that are typically only controlled by the administrator or domain owner (those listed in whois, postmaster@domain, etc.). This string has to be entered in the system (manually or by clicking a link) and then the user should confirm that he wants to add the domain name to the specified account

2.3.7.2. DNS TXT record — New Feature
The system generates a random string which needs to be put into a DNS TXT record for the domain. This record is then queried by the system

2.3.7.3. HTTP — New Feature
The system generates a random string which needs to be placed in a file on a web server running on that domain (URL should be specified by the system to avoid Pastebins and other sites allowing to host user content)

2.3.7.4. Whois — New Feature
The system generates a random string which has to be included in the whois database entry for that domain

2.3.7.5. Statement by 2 Assurers — New Feature
Statement by two Assurers about the ownership/control of this domain (only meant as last resort, maybe only under the control of a support engineer?)

New Feature
2.3.7.6. Other Methods
May be added over time. The system should allow for requiring more methods to verify the domain name if policy changes in the future

Existing Feature
2.3.8. Remove domain names
Certificates containing that domain name should be revoked, domain should only be marked as deleted until the last certificate expires

Altered Feature
2.3.9. Modify login method preferences (requirement 2.1.1.)
It should be verified whether the provided credentials work as described in requirement 1.8.2.2.

Existing Feature
2.3.10. Modify password recovery method preferences (requirement 2.2.)
A notification should be sent to the primary email address if critical information is displayed (e.g. when viewing common questions for password recovery ("maiden name of the mother" etc.))

Existing Feature
2.3.11. Modify announcement/newsletter settings

2.4. Web of trust

Existing Feature
2.4.1. List Assurances
And also the sum of points accumulated on each name. This page should also provide a summary of how to get more points and other requirements to move on. Traditionally there also has been some placement (e.g. you are the user with the 178th most Assurance points).

Existing Feature
2.4.2. Link to CATS
CATS (CAcert Assurer Testing System) is a system that offers multiple choice test on various subjects, most importantly the Assurer challenge which tests a prospective Assurer's knowledge about the Assurance process. The system itself is out of scope of the critical system.

Existing Feature
2.4.3. Import CATS results
The CATS regularly pushes the report about successful tests to the critical system. A passed Assurer challenge is required to become an Assurer.

Existing Feature
2.4.4. List absolved tests
The system should provide a list of passed tests so the user can verify if the test done on CATS has been successfully imported

Altered Feature
2.4.5. Link to the documentation
The explanation itself about how the WoT works is placed outside the scope of the critical system (e.g. CMS, wiki)

Altered Feature
2.4.6. Links to CAP-form generator
CAP forms are the forms used in the CAcert Assurance Process. The links should go to a pre-filled version (names, date of birth, primary email address encoded in the HTTP request parameters) and a plain version. The CAP form generator itself is placed outside the scope of the critical system.

Altered Feature
2.4.7. Link to the Assurer search
Link to the system that allows to find an Assurer in the same area. The system itself is outside the scope of the critical system (e.g. CMS)

New Feature
2.5. Trusted Third Party (TTP) programme
If the user has less than 100 Assurance points (precondition may change) there should be a link to the documentation for the TTP programme and a form to request a TTP Assurance. If used, a mail should be sent to the TTP Assurer team that includes TTP Assurances already present in the account.

2.6. Certificate Management

2.6.1. Requirements common to all certificate types

2.6.1.1. Certificate signing keys offline — Existing Feature
The keys for the subroot certificates may only be kept on machines not accessible from the network. Ideally they are additionally contained in a hardware cryptography device (e.g. a smartcard)

2.6.1.2. Only verified information may be included in the certificates — Existing Feature
Names may only be included if they have been assured with at least 50 Assurance points (the concrete preconditions may change). Email addresses and domain names may only be included if they have been verified (requirements 2.3.4. and 2.3.7.). Photo IDs as found in PGP keys and other information may not be signed.

2.6.1.3. Minimum key size — Existing Feature
The system should allow to globally specify minimum key sizes for some cryptographic algorithms. No certificate should be issued or renewed if the key doesn't fulfil that requirement

2.6.1.4. Key restrictions — Existing Feature
Apart from the key size other restrictions may be applied to the keys. The concrete restrictions should be modifiable.

2.6.1.5. Validity period — Existing Feature
Certificates are to be valid for six months for users with less than 50 Assurance points and 2 years for others, although the validity periods and point levels may change in the future.

2.6.1.6. Selection of issuing subroot — Existing Feature
The decision on which subroot will be used to issue the certificate is subject to change. With the current (as of January 2012) certificates there are only two root certificates: the so-called "class 1 root" which may be used for all certificates and the "class 3 root" only to be used by users who have more than or exactly 50 Assurance points (the user may choose to use the "class 1 root" nevertheless). With the new root structure this decision becomes more complex and the exact conditions are documented in the CPS (section §1.4.5)

2.6.1.7. List of issued certificates — Existing Feature
The system should provide a list of all certificates issued for the account. To improve orientation it should be possible to hide expired/revoked certificates and at least the following properties should be shown: common name, expiration date, revocation date, serial number, comment entered on creation, link to download

2.6.1.8. Download of the certificate — Existing Feature
A possibility to download the certificates should be offered. The download should be restricted to the user itself (and account sitters) for privacy reasons

2.6.1.9. Renewal of the certificate — Existing Feature
It should be possible to renew a certificate without going through the whole issuing process again (same properties and public key with another validity period)

2.6.2. Requirements common to X.509 certificates

Existing Feature

2.6.2.1. Submission via CSR (Certificate Signing Request)
The signature on the CSR should be checked and the information contained in it should be taken as default values for the certificate properties in the customisation and confirmation step

Existing Feature

2.6.2.2. Submission via web browser
The keys are interactively generated and the built-in methods to verify the possession of the private key (typically challenge/response based) are used. At least the key size should be selectable.

Altered Feature

2.6.2.3. Customisation and Confirmation of certificate properties
The user should confirm (e.g. because some properties in the CSR could not be fulfilled – maybe because of unverified domain names) and be able to edit the properties to be contained in the certificate (up to a certain point, e.g. subject alternative names may be edited, key usage may not)

Altered Feature

2.6.2.4. Revocation
The system should offer a way to revoke certificates (accessible from the certificate list). The revocation date should be set by the user or specify a date set to a date before the validity period. By no means always the current date should be used (if a signature was made before the revocation date it is sometimes considered still valid). Also the user should be able to select the reason why the certificate was revoked (to allow populating the reasons field in the CRL)

2.6.3. Requirements common to PGP certificates

Existing Feature

2.6.3.1. Submission
Because of the nature of PGP, a public key that already contains the properties wanted by the user is uploaded and the system can only check those properties for the concordance with the policies, they can not be edited. PGP keys can contain multiple identities and the system may selectively certify only those that can be verified (the comment field is ignored), this selection should be presented to the user in a confirmation step (including the reasons why the identities can not be certified if there are any), so that it's possible for the user to spot errors and correct them.

New Feature

2.6.3.2. Revocation
The user may upload a revocation certificate for each key which can be downloaded at any time by the user. This is basically an escrow service for revocation certificates. To prevent abuse as a file hosting platform and prevent errors, the revocation certificate should be checked for validity for the assigned PGP certificate

New Feature

2.6.3.3. Key server upload
In the future the ability to directly upload the signed key or on revocation a previously provided revocation certificate to a key server may be added

2.6.4. Certificate types
At least the following types of certificates need to be provided. Further types might be added in the future

New Feature

2.6.4.1. X.509 Login certificate
Does not include any identity information whatsoever and is therefore

not useful for third parties but may be used by the member to log in to the CAcert account using certificate authentication (mapping to the account via serial number). This is especially useful as the user can't create other types of client certificates until she is assured at least one time

2.6.4.2. X.509 Anonymous client certificate — Existing Feature
Only contains one or more verified email addresses but not a name ("CAcert WoT User" may be used instead)

2.6.4.3. X.509 Client certificate — Existing Feature
Contain a name (which has to match one of the assured names) and any number of verified email addresses

2.6.4.4. X.509 Server certificate — Existing Feature
Contain any number of verified domain names and subdomains thereof (no assured name of the member)

2.6.4.5. PGP certificate — Existing Feature
Contain any number of Assured name, verified email address pairs

2.7. Account Sitting — New Feature
Inexperienced users may give experienced users they trust limited access to their account to help them with certificate issuing, revocation, email/domain verification etc. without giving them their account credentials. Only assured users (i.e. users with more than 50 Assurance points, though this precondition may change) may act as account sitters. Account sitters may not access or adjust account login or recovery credentials

2.7.1. Explicit enabling
Users should need to explicitly add the account sitter in their account by entering the email address of the account sitter and a notification should be sent to the primary email address of the user including a link/token to confirm the change.

2.7.2. List of account sitters
The system should provide a list of persons, with assured name and email address, who have access to the account and provide a possibility to revoke access

2.7.3. Public key submission
The account sitter should be able to specify the properties of the desired certificate but not be able to submit the public key. The user should be provided with a possibility to submit the public key corresponding to the defined properties and get back the issued certificate.

2.8. Manage third party API permissions — New Feature
The user should be able to see all third parties that have access to his account (see requirement 11.) and what kind of services they are permitted to use. There should be a low-barrier possibility to revoke this access (full or in parts)

2.9. Excerpt of all personal data — New Feature
To comply with data protection regulations the user should be able to view all personal data CAcert has stored about him

3. Assurers
Assurer is a person with 100 or more Assurance points (at least one name assured to at least 100 points or at least one name assured by two TTP Assurances and

a TTP-TOPUP Assurance) and a passed Assurers challenge (these preconditions might change in the future). An Assurer should have all the functionality a normal user has (requirement 2.) and in addition to that the following:

3.1. Assure Someone
When entering the primary email address of a user the Assurer should be presented the following data and inputs:

Altered Feature

3.1.1. Reminder email
If the email address entered is not a primary email address of any account in the system, an option should be provided to send an invitation/reminder message to the email address. The language of the message should be selectable and an English version should always be included. If the email address is registered as additional email address of an account but not as primary email address instructions how to solve the problem should be given and the language set in the account should be used. In any case it should instruct the user to contact the Assurer once she has created the account or set the primary email address properly or the system should provide a way to notify the Assurer automatically once this has happened.

New Feature

3.1.2. Date of Birth
If the Assuree is less than 18 years old a special notice should be shown including a reference to the PoJAM (Policy on Junior Assurers/Members) should be given and the Assurer should have to explicitly mark that he did check the requirements given in there (usually parental consent). The exact age from which on this notice is triggered may change.

Existing Feature

3.1.3. Link to relevant policies/documents
That includes Assurance Handbook, Assurance Policy and Practice on Names. More might be added.

Altered Feature

3.1.4. Names
All names present in the Account should be shown to the Assurer with a possibility to assign each name a certain amount of Assurance points. This amount is limited to the maximum amount the Assurer may issue which is determined by how many Assurances the Assurer has already issued. Current policy is to start with 10 points that may be issued and then for every five Assurances five more points may be issued up to 35 points which is the maximum amount of points that may be issued by a single person. Additionally experience points earned by other means (Assurer Training Events etc.) are taken into account (2 experience points correspond to one Assurance given). This policy may however change in the future. If a name has already been assured by the Assurer the points may only be increased (i.e. if the name was previously assured with 10 points the Assurer may increase them up to the maximum amount he can currently issue but not 9 or less points). Directions should be given that this increase is only allowed if actually another Assurance did take place, not to correct errors or make use of gained experience points. Assurances to the same user only count once towards the experience level. It should be possible to add functionality to specify in which ID documents the name is contained and show that information to future Assurers (this is to allow detecting faulty Assurances done by other Assurers).

Existing Feature

3.1.5. Location
The location where the Assurance took place (free text)

3.1.6. Date — Existing Feature
The date on which the Assurance took place. This may be a validated field but the system should respect that in old data from the existing system a free field was used.

3.1.7. Agreements the Assuree entered into — Altered Feature
There are some Agreements that are captured during the Assurance process (currently only the CCA). During the introduction of such Agreements into the process there usually are some Assurances which do not include those Agreements yet (or old Assurances made before the introduction need to be entered). If a particular currently required Agreement is not marked as being entered into by the Assuree, the system should ask the Assurer, whether this was accidental or if indeed the Assuree did not enter into that Agreement. Care should be taken that this question does not lead Assurers to mark the Agreement as entered into while in reality it was not. The set of Agreements the Assuree may enter into as well as the subset of those that currently are required may change.

3.1.8. Out of band password for requirement 2.2.1. — New Feature
Filling this field is optional

3.1.9. Statement by the Assurer — Existing Feature
The Assurer explicitly needs to mark that he has met the Assuree in person and that it was conducted according to the Assurance Policy (for legal reasons) including a link to the relevant policies.

3.1.10. Confirmation/Notification Email — Existing Feature
Once an Assurance has been successfully entered a notification email should be send to the Assurer and the Assuree. The mail to the Assurer should contain the points assigned to each name but not the email address of the Assuree (to avoid Spammers collecting many addresses if they are able to gain access to an email account of a very active Assurer). The email to the Assuree should contain the newly gained points as well as the total amount for each assured name and the status of each name (assured/not assured) and if the user has enough points to be an Assurer but not passed the Assurer challenge yet there should be a pointer to that.

3.1.11. Optimisation for multiple Assurances — Existing Feature
The Assure someone functionality should be optimised for entering multiple Assurances in a row which is often the case for Assurers who are present at open source conventions and similar events. For example date and location should be pre-filled with the data from the previous Assurance entered in that session.

3.2. Link to register as listed Assurer — Altered Feature
The registration as listed Assurer itself is outside the scope of the critical system (e.g. CMS) but there should be a prominent link to get registered.

3.3. List of issued Assurances — Altered Feature
Similar to requirement 2.4.1. the system should provide a list of all Assurances the Assurer has entered, how many points the Assurer may issue and the position of the Assurer in the ranking. If an Assurance was entered less than 24 hours ago, a possibility should be provided to send a message requesting the revocation of the Assurance to support including the identifier of the Assurance, an authorisation token (see requirement 10.2.) and a reason why the Assurance should be revoked.

New Feature
3.4. Request code signing
An Assurer may request to enable code signing certificates for his account. This request triggers an email to support including an authorisation token (requirement 10.2.) which then grants or denies the request. In the future this may be automated or enabled by default.

Existing Feature
3.5. X.509 Code signing certificates
An Assurer who has code signing enabled may in addition to the other certificate types get a code signing certificate. Which is basically a client certificate with a special extension. This may be realised by adding another check box to the customisation and confirmation step (requirement 2.6.2.3.).

4. Organisation Admins
Registered organisations may apply for an organisation account. This is actually not an extended user account but the organisation may name several organisation admins which then get the following additional functionality on their normal user account.

Existing Feature
4.1. Add another organisation admin
Only users who already are Assurers may be added as organisation admin (this precondition might be adjusted in the future).

4.1.1. Email
To identify the user account to promote to organisation admin

4.1.2. Department
The department the new admin is working in

4.1.3. Comment
An optional comment (e.g. why the user has the admin permission) the use of which is free to define by the organisation admins

New Feature
4.1.4. Notifications
All other admins of the organisation should be notified of the change

Existing Feature
4.2. Remove another organisation admin

New Feature
4.3. Request addition/removal of a domain name
On request an email should be sent to the Organisation Assurers who then do the verification and action

Existing Feature
4.4. Organisation certificates
Organisation certificates are X.509 certificates (client, server, code signing) that also contain the organisations details (name, state, country etc.) as specified in requirement 5.2. Organisation client and code signing certificates are handled a little bit different than in normal user accounts: the email addresses do not need to be verified separately, any email address where the domain part is a domain name of the organisation (or subdomain thereof) is accepted. Also the name is not verified, the organisation admin may just put in any name he sees fit. Organisation client certificates can not be used to log into CAcert accounts

New Feature
4.5. Automation API
The system should offer an API that allows organisation administrators to automate their issuing process. This API should use special authorisation credentials (e.g. tokens or specially marked certificates) that can only be used for accessing the API and not the full account. The system should also offer a way to revoke those credentials and view the last actions performed with those credentials

5. Organisation Assurers — Existing Feature
   Organisations are verified and managed by Organisation Assurers. Organisation Assurers are users who have the "Organisation Assurer" role set. Collaboration between Organisation Assurers is done via an external system connected via email (e.g. issue tracker or mailing list). Following functionality is available in addition to the features for normal users:

   5.1. List organisations
   List all organisations in the system. There should be means to filter, sort and search this list as it might get quite huge.

   5.2. Add organisation
   Once the Organisation Assurer has verified the organisation he adds the organisation to the system. The following information has to be included: Name, contact email , town, state, country (ISO 3166-1 alpha-2 code) and a comment (the comment field has to be big enough to comfortably enter multiple sentences)

   5.3. Edit organisation

   5.4. Delete organisation

   5.5. Manage Organisation Admins
   Same as requirements 4.1. and 4.2.

   5.6. Add domain name
   The ownership will be verified manually by the Organisation Assurer. If the domain name already exists in the system an error should be given.

   5.7. Remove domain name
   Similar to requirement 2.3.8.

6. Trusted Third Party (TTP) Assurer
   The information gathered via TTPs is verified by TTP Assurers. TTP Assurers are users who have the "TTP Assurer" role set. The following functionality is available in addition to the features for normal users:

   6.1. List of TTP Assurances — New Feature
   On entering the primary email address of a user the TTP Assurer should be able to see a list of names and corresponding TTP Assurances already present on the account.

   6.2. Enter TTP Assurance — Altered Feature
   On entering the primary email address of a user, the TTP Assurer should be able to add another TTP Assurance. If there are already two TTP Assurances present on the account an error message should be given instead. Similar to requirement 3.1. the date of birth and a list of names should be shown. An TTP Assurance always gives 35 points to the names that could be verified (concrete amount might change in the future). Similar to requirement 3.1. a list of ID documents which contain the name could be added in the future. Additionally the country, name, location and registration number of the TTP, date of the TTP verification taking place, Agreements the user entered into should be recorded.

7. TTP-TOPUP Assurer — New Feature
   Via TTP Assurances a user can only gain 70 Assurance points. TTP-TOPUP is a

programme to educate the user about CAcert's Assurance process so he can become a CAcert Assurer. This is done by TTP-TOPUP Assurers which are TTP Assurers with the "TTP-TOPUP Assurer" role set. The following functionality is available in addition to the features for TTP Assurers:

7.1. Enter TTP-TOPUP Assurance
In contrast to normal Assurances and TTP Assurances the TTP-TOPUP Assurance is not a statement about the identity of the user but the education. Therefore it should apply to the whole account not only single names and does not count towards the verification of names. It only counts towards the Assurer capability. On entering the primary email address of a user, the TTP-TOPUP Assurer may mark the account as TTP-TOPUP assured.

New Feature

8. Arbitrator
Arbitrators handle difficult cases that are not covered by previously described and authorised procedures. They are the "judiciary" of CAcert. Arbitrators may give authorisation for certain actions executed by the Support Engineer. Arbitrators are users who have the "Arbitrator" role set. The following functionality is available in addition to the features for normal users:

8.1. Issue authorisation tokens
On entering the primary email address of an account and an Arbitration case number, the Arbitrator may select one or more actions to be authorised and get back a token that he can send to support (see requirement 10.2.).

New Feature

8.2. Manage precedence cases
The Arbitrator might define a precedence case which has some authorised actions associated with it like a token, but that can be used multiple times provided the executing Support Engineer records the support case number. There also should be a list of existing precedence cases with a possibility to revoke/disable them and view the list of recorded support cases

9. Board Member
Board members are users elected by the community. For the system Board members are users who have the "Board" role set. The following functionality is available in addition to the features for normal users:

New Feature

9.1. Issue authorisation tokens to change roles
On entering the primary email address of an account and a Board motion number, the Board member may select one or more roles assigned or removed and get back a token that he can send to support (see requirement 10.2.).

Existing Feature

9.2. See all privileged users
In the privileged roles review (requirement 12.2.) board members always see all privileged users not only those with the same role.

10. Support Engineers
Support Engineers answer help requests by users, execute Arbitration rulings and execute many daily work tasks. They are part of the "executive" of CAcert. Support Engineers are users who have the "Support Engineer" role set. The following functionality is available in addition to the features for normal users:

Existing Feature

10.1. View basic details about an account
On entering an email address of an account or the account ID at least the following details should be provided but it should be possible to easily add more in the future:

10.1.1. Names
Show all names and the Assurance points assigned to them

10.1.2. Date of birth

10.1.3. Account status
Whether the account has been blocked or blocked from making Assurances

10.1.4. Amount of experience points

10.1.5. Email addresses
The primary email address should be marked as such

10.1.6. Domain names

10.1.7. Privileges
Such as assigned roles, if the user is an Assurer or not, whether code signing is enabled etc.

10.1.8. Passed Tests
Such as Assurers Challenge, Triage Challenge etc.

10.1.9. Account activity
A rough guidance about whether the account is still in active used (last login within last month/year etc.) and when the account was initially created

10.1.10. Overview over issued certificates
How many certificates have been issued? How many of those have been revoked/expired/are still valid? What is the date when the last certificate expires (even if it has been revoked)?

10.1.11. Debugging information
Show technical anomalies in the data records about the user if they are present

10.2. Privileged functionality — New Feature
This functionality is only available when a token (a random string that is generated when a privileged action is authorised) is entered or a precedence case is selected and a support case number is entered. It should be easily possible to add more functionality in the future

10.2.1. Read only view of all user details
This should be partitioned by functionality

10.2.2. Edit user data — Altered Feature
See requirement 2.3.. Allow date of birth to be edited even if there are Assurances on some name

10.2.3. Revoke Assurances — Existing Feature
Allow to remove an Assurance. The date when the Assurance was entered should be displayed

10.2.4. Revoke certificates

10.2.5. Remove account sitters

10.2.6. Remove API permissions

10.2.7. Assign or remove roles

11. Third party API — New Feature
The system should offer an API that provides some services to third parties

11.1. Only approved parties
The API should only offer the services to registered and approved parties. This approval should also contain the subset of the services the party is generally allowed to use

11.2. Only on users permission
If the user has not granted permanent permission (see requirement 2.8.) he should be asked if the third party may access his data. This request should contain a list of services the party wants to use with the ability to selectively deny the access and the possibility to choose whether to allow this access once or permanently (until revoked)

11.3. Provided services
At least the following services should be provided, others should be easy to add in the future:

11.3.1. Name verification
Given an email address and a name, the system responds whether the data matches. It should only take assured names into account (i.e. names that have at least 50 Assurance points; concrete precondition might change)

11.3.2. Date of birth verification
Given an email address and a date of birth or age the system responds whether the data matches. Only assured accounts are to be taken into account (i.e. one name assured to at least 50 Assurance points; concrete precondition may change)

11.3.3. Age verification
Given an email address, a country and a purpose (e.g. old enough to be legally competent, view content not suited for minors, buy alcoholic beverages) the system should respond whether the user is according to the rules in that country allowed for that purpose (only taking age and not other conditions into account). The system should give an error if the age for the country is not known or does not apply. Only assured accounts are to be taken into account (i.e. one name assured to at least 50 Assurance points; concrete precondition may change)

11.3.4. Assurer status
Given an email address the system responds whether the user has Assurer status

11.3.5. Assurance points that may be issued
Given an email address the system responds with the amount of Assurance points the user may issue (0 if the user is not an Assurer yet, see requirement 3.1.)

12. Additional Requirements

New Feature
12.1. Expiration of Assurance and/or Experience points
The system should allow for changing the point system in such a way, that points expire after some fixed time span or only have an effect if the last Assurance is less than a fixed time span in the past (e.g. a non-anonymous certificate may only be issued if the last Assurance was less than 2 years ago). This might be required by changing external demands.

Altered Feature
12.2. Review of privileged roles
All users who have a privileged role should have access to a list containing all

other users who also have that same privileged role. Also a regular permission report in the form of an email containing the same information should be sent to those users.

A.2. Non-Functional Requirements

1. Security
   As a CA security is a major concern. A detailed list of valuable assets and the protection required can be found in the next section
2. Audit trail
   All critical operations should be recorded so that in the case of an incident an investigation of the cause can be conducted.
3. Scalability
   Once the Audit is complete we expect a substantial rise in users. Therefore it should be easily possible to scale the system by adding further hardware or other easy-to-deploy modifications
4. Maintainability
   In the security business often new requirements show up. It should be easy to extend the system to adjust to these changing requirements. Major points where change is already to be expected are mentioned in the functional requirements but other points may also be affected
5. Testability
   The system should allow for thorough testing. Testing is employed to ensure the quality of new code and prevent the introduction of regressions
6. Protection of personal data
   The protection of the personal data of the users is a major concern to CAcert. Also it is a requirement imposed by legal regulations.

   6.1. Data economy
   Personal data should only be used if needed. And if used it should be used in a way that minimises the impact.
7. Usability
   In order to facilitate widespread use, the system should be usable by persons of average technical skills.

A.3. Critical Assets

In this section assets that need special protection are identified along with the kind of protection required and the reason for this protection and the impact if it fails.

A.3.1. Web of Trust

A.3.1.1. Integrity

If the integrity of the web of trust was compromised, an attacker could fake the verification of his identity and therefore create certificates with unverified information in them without having to fear prosecution by backtracking through the web of trust. Changes in the web of trust should be recorded so that after-the-fact investigations if the integrity might have been compromised are possible. If the integrity of the web of trust is compromised and it is not possible to either investigate and recover from the damage or setting the system back to a known safe state this would mean that CAcert would have to rebuild the web of trust from scratch and therefore probably mean the end of existence for the project.

A.3.1.2. Confidentiality

To protect information about relationships between users, which might be considered personal data, the connections in the web of trust should not be public. If the relationship data leaks it might result in bad publicity.

A.3.2. Login/Recovery Credentials

A.3.2.1. Integrity

If an attacker is able to set new credentials he has full access to the account and can issue certificates in that users name and if the user has Assurer status he might falsely verify his own account for a fake identity. If multiple accounts get compromised in this way (especially Assurer accounts) and it can be determined to which accounts this applies, a whole subgraph might need to be removed from the web of trust which is not desirable but feasible. If it is not possible to determine which accounts have been compromised then essentially the integrity of the web of trust is compromised and the same ways of recovery or failure to do so apply. Changes of login or recovery credentials should be recorded so that an investigation of accounts which might have been compromised is possible.

A.3.2.2. Confidentiality

A user might against all advice use the same password for multiple services so even if CAcert is compromised the credentials stored should not be useful for logging into other services. If an attacker recovers the credentials for an account in a usable way he has full access to the account and might perform any attack described in the previous section. Even if the credentials are recovered in an unusable way it might still result in bad publicity.

A.3.3. User Data

A.3.3.1. Integrity

If an attacker is able to modify the data of his account after he already got Assured or add an email address or domain he does not have control over, he might issue certificates for a fake identity. These fake identity certificates might in turn be used to compromise other systems (e.g. secure network connections) and result in a major loss of trust in CAcert certificates. Changes of user data should be recorded so that an investigation of unauthorised changes is possible.

A.3.3.2. Confidentiality

To protect personal data about the user, this information should not be public. Assurers might access some of this data on entering the primary email address in order to assure the user. Failure to enforce confidentiality of private user data might result in bad publicity and many users requesting to delete their data.

A.3.4. Issued Certificates

A.3.4.1. Integrity

While integrity of the certificates themselves stored for the retrieval by the user is not required as they contain inherent integrity checks, some metadata, like which certificate belongs to which user account is critical. If an attacker could associate his certificate with the user account of another user he might be able to use it to login as that user (especially for login-only certificates), effectively resulting in an integrity violation of login credentials. All issuing of certificates should be recorded so that false metadata can be identified.

A.3.4.2. Confidentiality

Certificates should not be public as they might contain personal user data. If the confidentiality is violated it might result in bad publicity for CAcert.

A.3.5. Revocation Information

A.3.5.1. Integrity

If an attacker is able to mark a compromised certificate as not revoked, he still may use this certificate even if the software validating the certificate does revocation checking. Therefore every revocation of a certificate should be recorded so that an investigation and recovery is possible. If such an attack succeeds it might have an impact on trust in CAcert certificates.

A.3.5.2. Availability

If the revocation status of a certificate is not available at the time of verification, the software verifying might react by either warning the user or even marking the certificate invalid or just ignoring this kind of error. Warnings or treating the certificate as invalid might give the impression that the connection is insecure when in reality this is not the case or even worse ignoring the error might lead to putting trust in a certificate that has been compromised. Low availability of revocation information gives the impression that the CAcert service is unreliable, web site owners might complain about inaccessibility for their users and software vendors will ignore revocation status checking errors because of the infeasibility of strict checking. The revocation status providing services should be redundant for these reasons.

A.3.6. Root/Subroot Certificates

A.3.6.1. Integrity

If an attacker is able to replace the root and subroot certificates with his own certificates, users coming to the CAcert website to install the CAcert root certificates might install the attackers certificate. Although there also is the unsolvable problem that the data is changed in transit, this would probably only affect a few connections while replacing the certificates on the website would be persistent and therefore affect more users. The real solution to the problem would be that users verify the fingerprint of the certificate which is distributed out-of-band or receive the certificate in another secure way, but to be realistic one has to assume that this is not done by every user. If an attacker is able to replace the root certificate this might result in bad publicity for CAcert.

A.3.7. Certificate Signing Keys

A.3.7.1. Integrity

If an attacker is able to modify the keys used to sign end user certificates it might pose some interruptions in our service as highly secured backups would have to be restored.

A.3.7.2. Confidentiality

If an attacker is able to retrieve the certificate signing keys or execute some operations on it (e.g. signing arbitrary values) this would enable him to sign arbitrary certificates (under the current root setup he might even issue a subroot he might use for himself). A way to recover from this failure might be to revoke the compromised certificate signing keys and generate new ones. All signing operations should be logged to allow investigation and

strict input validation is required to hinder signing of dangerous values (e.g. a blacklist for critical domain names that may never be contained). Nevertheless if the confidentiality is violated it would result in a major loss of trust in CAcert certificates that would threaten the existence of the project.

B. Layered Design

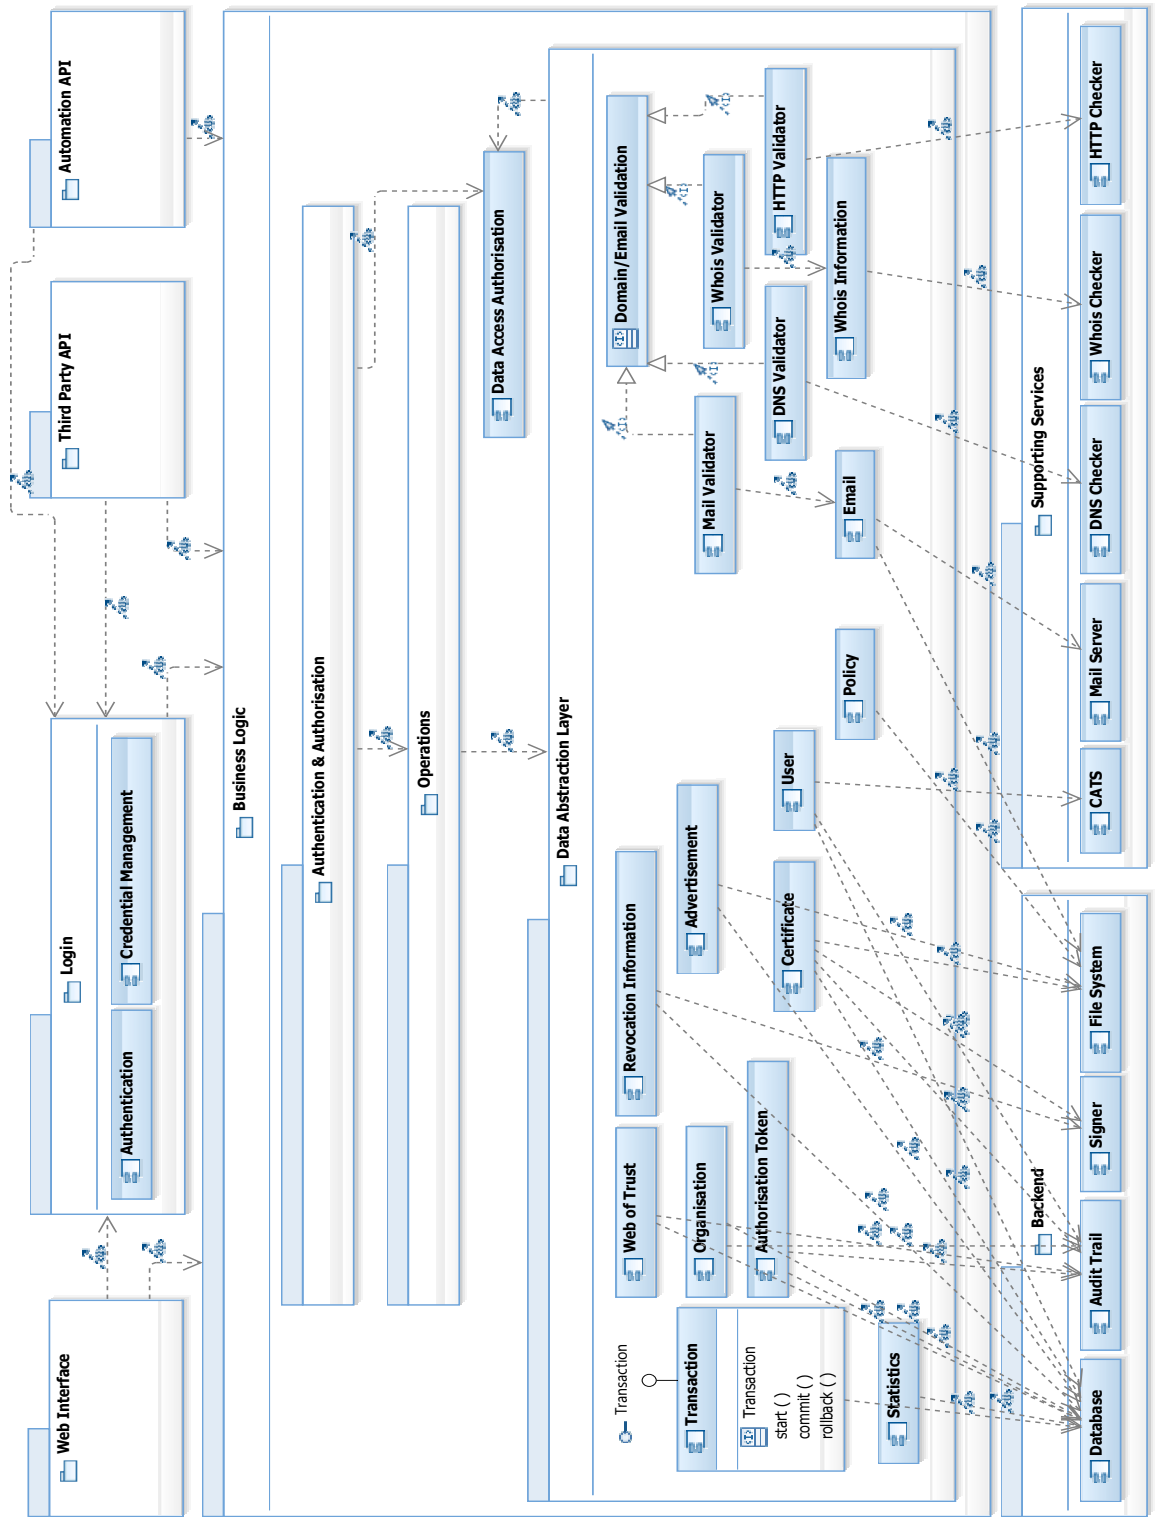

C. Layered Design Analysis

	Effort/Risk/ Gain
C.1. Web of Trust	
C.1.1. Unauthorised Modification of the Web of Trust	
1. Direct access to the database	**2/1/8** **5/2/7**
1.1. Recover database credentials AND	**5/1/3**
1.1.1. Recover application credentials	**5/1/3**
1.1.1.1. Exploit information leak (e.g. via error messages, stack traces and other debugging information)	**5/2/2**
1.1.1.2. Exploit memory access violation (buffer overflow attack, printf attack etc.)	**6/1/8**
1.1.1.3. File content leak of configuration files (user controlled paths, unintentional check-ins into version control etc.)	**5/1/3**
1.1.1.4. Intercept credentials from SQL client connection	**8/2/5**
1.1.1.4.1. Intercept connection on the network AND	**4/2/4**
1.1.1.4.1.1. Punch hole through firewall	**6/4/5**
1.1.1.4.1.2. Attack that gains user privileges on hosts on the internal network	**4/2/5**
1.1.1.4.2. Attack encryption protocol	**8/1/9**
1.1.1.5. Set new credentials for application (e.g. via SQL injection)	**6/2/7**
1.1.1.6. Social Engineering	**5/6/6**
1.1.2. Recover administrative credentials	**6/7/8**
1.1.2.1. Social Engineering	**6/7/8**
1.1.2.2. Attack that gains root privileges on the system running the database (exploit of security vulnerability for the OS or other component running with root privileges, access to hardware etc.)	**7/2/8**
1.2. Connect to the database	**4/2/3**
1.2.1. Punch hole through the firewall and connection exposed to internal hosts	**6/4/5**
1.2.2. Attack that gains user privileges on hosts on the internal network and connection exposed to internal hosts	**4/2/5**
1.2.3. Attack that gains user privileges on the database host	**4/2/5**
2. Intercept and modify commands sent to the database	**8/2/7**
2.1. Database connection exposed to internal hosts AND	
2.2. Intercept and modify traffic on the internal network AND	**4/2/4**

6/4/5 2.2.1. Punch hole through firewall

4/2/5 2.2.2. Attack that gains user privileges on hosts on the internal network

7/2/8 2.2.3. Attack that gains root privileges on the database host

8/1/9 2.3. Attack encryption protocol

6/2/7 3. Inject malicious database commands into the commands sent by the data object (SQL injection)

3/2/3 3.1. Bypass input validations at upper layers (e.g. by using unusual escape sequences that will be unescaped at lower layers) AND

6/1/4 3.2. Execute the injected data in the context of SQL commands (e.g. improperly handled variable parts in SQL templates)

2/1/7 4. Modify data objects at the data abstraction layer

7/1/7 4.1. Directly modify data objects

6/1/4 4.1.1. Bypass operations layer (buffer overflow, etc.) AND

5/1/4 4.1.2. Pass data access authorisation checking

5/1/4 4.1.2.1. Find a sequence of primitive operations leading to the intended state that are allowed for the attacker

6/1/4 4.1.2.2. Use authentication state of another user with enough privileges who is currently logged in (e.g. memory access violation)

6/1/5 4.1.2.3. Bypass the checking of authentication state (e.g. buffer overflow)

6/1/4 4.2. Make an operation the attacker is authorised to execute, modify the data objects in unintended ways

5/1/4 4.2.1. Pass data access authorisation checking (see subtree 4.1.2.) AND

4/1/3 4.2.2. Exploit operation logic (e.g. sloppy condition checking, buffer overflows)

2/1/6 4.3. Execute an operation the attacker is not allowed to execute

7/1/6 4.3.1. Bypass authorisation checking

6/1/5 4.3.1.1. Bypass authorisation checking on the authentication & authorisation layer (exploit sloppy error checking, buffer overflows etc.) AND

5/1/4 4.3.1.2. Pass data access authorisation checking (see subtree 4.1.2.)

7/1/6 4.3.2. Cause the authentication & authorisation layer to call a different operation than the one that has been checked for authorisation

6/1/5 4.3.2.1. Exploit authentication logic (e.g. sloppy condition checking, buffer overflows) AND

5/1/4 4.3.2.2. Pass data access authorisation checking (see subtree 4.1.2.)

2/1/6 4.3.3. Use the authentication of another user with the required privileges

2/3/4 4.3.3.1. Social engineering/cooperation by the user (e.g. phishing)

2/1/6 4.3.3.2. Learn authentication credentials of the user (see section C.2.2)

2/1/6 4.3.3.3. Set new authentication credentials for the user (see section C.2.1)

4/2/5 4.3.3.4. Hijack authentication state of the user

4.3.3.4.1. Learn reauthentication credentials (e.g. cookies – analogous to C.2.2) AND **2/1/4**

4.3.3.4.2. Pass reauthentication checks, such as IP address restrictions **4/2/4**

4.3.3.5. Attack the authentication or reauthentication method (e.g. find a way to successfully authenticate without knowing the password) **6/1/7**

C.1.2. Violation of the Confidentiality of the Web of Trust

2/1/5

1. Direct access to database (see section C.1.1 subtree 1.) **5/2/7**

2. Intercept data received from the database **7/2/7**

2.1. Database connection exposed to internal hosts AND

2.2. Intercept traffic on the internal network AND **4/2/4**

2.2.1. Punch hole through firewall **6/4/5**

2.2.2. Attack that gains user privileges on hosts on the internal network **4/2/5**

2.2.3. Attack that gains root privileges on the database host **7/2/8**

2.3. Attack encryption protocol if used **8/1/9**

3. Read data from the audit log **6/2/6**

3.1. Get read access on the internal network **6/2/6**

3.1.1. Connect to the audit log service (see section C.1.1 subtree 1.2.) AND **4/2/3**

3.1.2. Exploit audit log service to read from the log instead of appending to it **6/1/3**

3.2. Get read access on the audit trail host **6/2/6**

3.2.1. Get administrative access on the audit trail host (see section C.1.1 subtree 1.1.2.2.) **7/2/7**

3.2.2. Read audit log with user access on the audit trail host **6/2/6**

3.2.2.1. Attack that gains user privileges on the audit trail host AND **4/2/3**

3.2.2.2. Attack that gains read privileges on the audit log file to non-administrative users on the audit trail host (via file content leak or user controlled paths in programmes running as super user etc.) **6/1/4**

4. Inject malicious database commands into the commands sent by the data object (SQL injection, see section C.1.1 subtree 3.) **6/2/7**

5. Access data through the data abstraction layer **2/1/6**

5.1. Directly access data objects (see section C.1.1 subtree 4.1.) **7/1/7**

5.2. Make an operation the attacker is authorised to execute, unintentionally disclose data (see section C.1.1 subtree 4.2.) **6/1/4**

5.3. Execute an operation the attacker is not allowed to execute (see section C.1.1 subtree 4.3.) **2/1/6**

C.2. Login Credentials

C.2.1. Unauthorised Modification of the Login Credentials

2/1/6
5/2/7 1. Direct access to the database (see section C.1.1 subtree 1.)

8/2/7 2. Intercept and modify commands sent to the database (see section C.1.1 subtree 2.)

3. Inject malicious database commands into the commands sent by the data object (see
6/2/7 section C.1.1 subtree 3.)

2/1/6 4. Modify credentials at the data abstraction layer

7/1/7 4.1. Directly modify data objects (see section C.1.1 subtree 4.1.)

4.2. Make an operation the attacker is authorised to execute, modify the data objects in
6/1/4 unintended ways (see section C.1.1 subtree 4.2.)

4.3. Execute an operation the attacker is not allowed to execute (e.g. set new login
2/1/6 credentials)

7/1/6 4.3.1. Bypass authorisation checking (see section C.1.1 subtree 4.3.1.)

4.3.2. Cause the authentication & authorisation layer to call a different operation than the one that has been checked for authorisation (see section C.1.1 subtree 4.3.2.)
7/1/6

2/1/6 4.3.3. Use the authentication of another user with the required privileges

2/3/4 4.3.3.1. Social engineering/cooperation by the user (e.g. phishing)

2/1/6 4.3.3.2. Learn (old) authentication credentials of the user (see section C.2.2)

4.3.3.3. Pass account recovery methods (e.g. intercept confirmation email sent to the primary email address and learn enough about the user to answer the reset
3/2/6 questions)

4/2/5 4.3.3.4. Hijack authentication state of the user (see section C.1.1 subtree 4.3.3.4.)

4.3.3.5. Attack the authentication or reauthentication method (e.g. find a way to
6/1/7 successfully authenticate without knowing the password)

C.2.2. Violation of the Confidentiality of the Login Credentials

2/1/7
1. Intercept credentials on the machine of the user (trojan, key logger, cross-site-scripting
3/2/6 etc.)

2. Deceive the user to believe a server controlled by the attacker is the server of CAcert
2/1/6

2.1. Make the user connect to a server the attacker controls, instead of the real server, and relay the authentication protocol to the real server (phishing, DNS attacks etc.)
1/1/4 AND

2.2. Bypass server side authentication of the encryption protocol (slightly different spelling in the domain name, attacks against domain name comparison like null characters,
2/1/6 use no authentication at all etc.)

3. Intercept credentials on the wire (not possible for challenge/response authentication
8/2/6 protocols)

2/2/3 3.1. Intercept the connection between the user and the server AND

3.2. Break confidentiality of encrypted connection **8/1/8**

4. Relay attack on challenge/response protocol **8/2/6**

4.1. Route connection through a host the attacker controls (e.g. by attacking name resolution, routing protocols or social engineering techniques) AND **2/2/4**

4.2. Break integrity and confidentiality of encrypted connection **8/1/9**

5. Intercept credentials on the server **5/1/7**

5.1. Attack that gains administrative privileges on the login server (similar to section C.1.1 subtree 1.1.2.2.) **6/2/8**

5.2. Exploit the login server to store credentials and expose them on successive requests **5/1/7**

5.2.1. Persist credentials (e.g. provoke an error that puts the credential as debug information in a log file) AND **3/1/5**

5.2.2. Restore persisted credentials (e.g. file content leak on a log file) **5/1/6**

5.3. Exploit the login server to communicate login credentials to the attacker (e.g. via code execution attack) **5/1/7**

C.3. User Data

C.3.1. Unauthorized Modification of User Data

Analogous to the modification of the web of trust data (section C.1.1). **2/1/8**

C.3.2. Unauthorized Access to User Data

Analogous to exposing data from the web of trust (section C.1.2). **2/1/5**

C.4. Issued Certificates

C.4.1. Modification of Issued Certificates

5/1/7

1. Direct access to the database with modify access to the certificate information **6/7/8**

1.1. Recover administrative database credentials (see section C.1.1 subtree 1.1.2.) AND **6/7/8**

1.2. Connect to the database (see section C.1.1 subtree 1.2.) **4/2/3**

2. Make the data object report false metadata (e.g. the user account the certificate is associated with or whether it may be used for login, by using memory violations, sloppy error checking etc.) **5/1/6**

C.4.2. Unauthorized Access to Issued Certificates

Analogous to exposing data from the web of trust (section C.1.2). **2/1/5**

C.5. Revocation Information

C.5.1. Unauthorized Modification of Revocation Information

6/7/6

1. Direct access to the database with modify access to the certificate revocation information

6/7/8

1.1. Recover administrative database credentials (see section C.1.1 subtree 1.1.2.) AND

6/7/8

4/2/3 1.2. Connect to the database (see section C.1.1 subtree 1.2.)

2. Modify the data at the point of distribution into the system of OCSP responders and

8/2/6 CRL caches

7/2/6 2.1. Forge the signature on the OCSP response or CRL AND

7/2/8 2.1.1. Get access to the certificate signing keys (see section C.7.2)

9/1/9 2.1.2. Forge the cryptographic signature

8/2/7 2.2. Modify the data on the wire (man in the middle attack)

2.2.1. Route the connection from the OCSP responder or CRL cache to the authoritative revocation server to a host the attacker controls (e.g. by attacking name

2/2/4 resolution, routing protocols or social engineering techniques) AND

8/1/8 2.2.2. Break the authenticity of the encrypted connection

7/2/6 3. Modify the data in the system of OCSP responders and CRL caches

7/2/6 3.1. Forge the signature on the OCSP response or CRL (see subtree 2.1.) AND

4/2/6 3.2. Get write access to an OCSP responder or CRL cache

4. Modify the data on the wire between the client and the system of OCSP responders

7/2/6 and CRL caches

7/2/6 4.1. Forge the signature on the OCSP response or CRL (see subtree 2.1.) AND

4.2. Route the connection from the client to an OCSP responder or CRL cache to a host the attacker controls (e.g. by attacking name resolution, routing protocols or social

2/2/6 engineering techniques)

C.5.2. Prevent Access to Revocation Information

5/3/5

1. Route a significant number of connections from clients to OCSP servers or CRL caches to a host the attacker controls or a dead end (e.g. by attacking name resolution, routing

5/5/6 protocols or social engineering techniques)

2. Denial of service attack against each of the distributed OCSP responders and CRL

6/3/5 caches

3. Denial of service attack against the point of distribution into the system of OCSP responders and CRL caches (single point of failure but possible to use white list filtering

5/3/3 in the router)

4. Denial of service attack against the signing server (e.g. by requesting and revoking

3/8/7 many certificates at once)

C.6. Root/Subroot Certificates

C.6.1. Modification of Root/Subroot Certificates

6/2/6

1. Write access to the file system on the front end server **6/2/8**

1.1. Attack that gains root privileges on the front end server (similar to section C.1.1 subtree 1.1.2.2.) **6/2/8**

1.2. Exploit some component or service running with elevated privileges on the front end server to overwrite files **6/2/8**

C.7. Certificate Signing Keys

C.7.1. Modification of the Certificate Signing Keys

8/2/6

1. Exploit some component or service running with elevated privileges on the signing server to overwrite the certificate signing keys AND **8/2/6**

2. Access the serial connection (see section C.7.2 subtree 1.2.) **6/2/8**

C.7.2. Violation of the Confidentiality of the Certificate Signing Keys

7/2/8

1. Exploit the signing server via the serial access protocol **7/2/8**

1.1. Exploit the signing server to expose information about the key material (e.g. by causing part of the signing key material to be included in a signed certificate via memory violation) AND **7/2/8**

1.2. Access the serial connection **6/2/8**

1.2.1. Directly access the serial connection (e.g. by getting root access on the host running the data abstraction layer, see section C.1.1 subtree 1.1.2.2.) **7/2/8**

1.2.2. Route malicious data through the normal signing process (similar to the SQL injection, see section C.1.1 subtree 3.) **6/2/7**

2. Gather information about the signing key material by observing side channels **7/5/7**

2.1. Discover a side channel (timings, power consumption, cache access patterns etc.) AND **5/2/4**

2.2. Access the side channel (e.g. by analysing response times or measuring power consumption) **7/5/4**

D. Service-Oriented Design

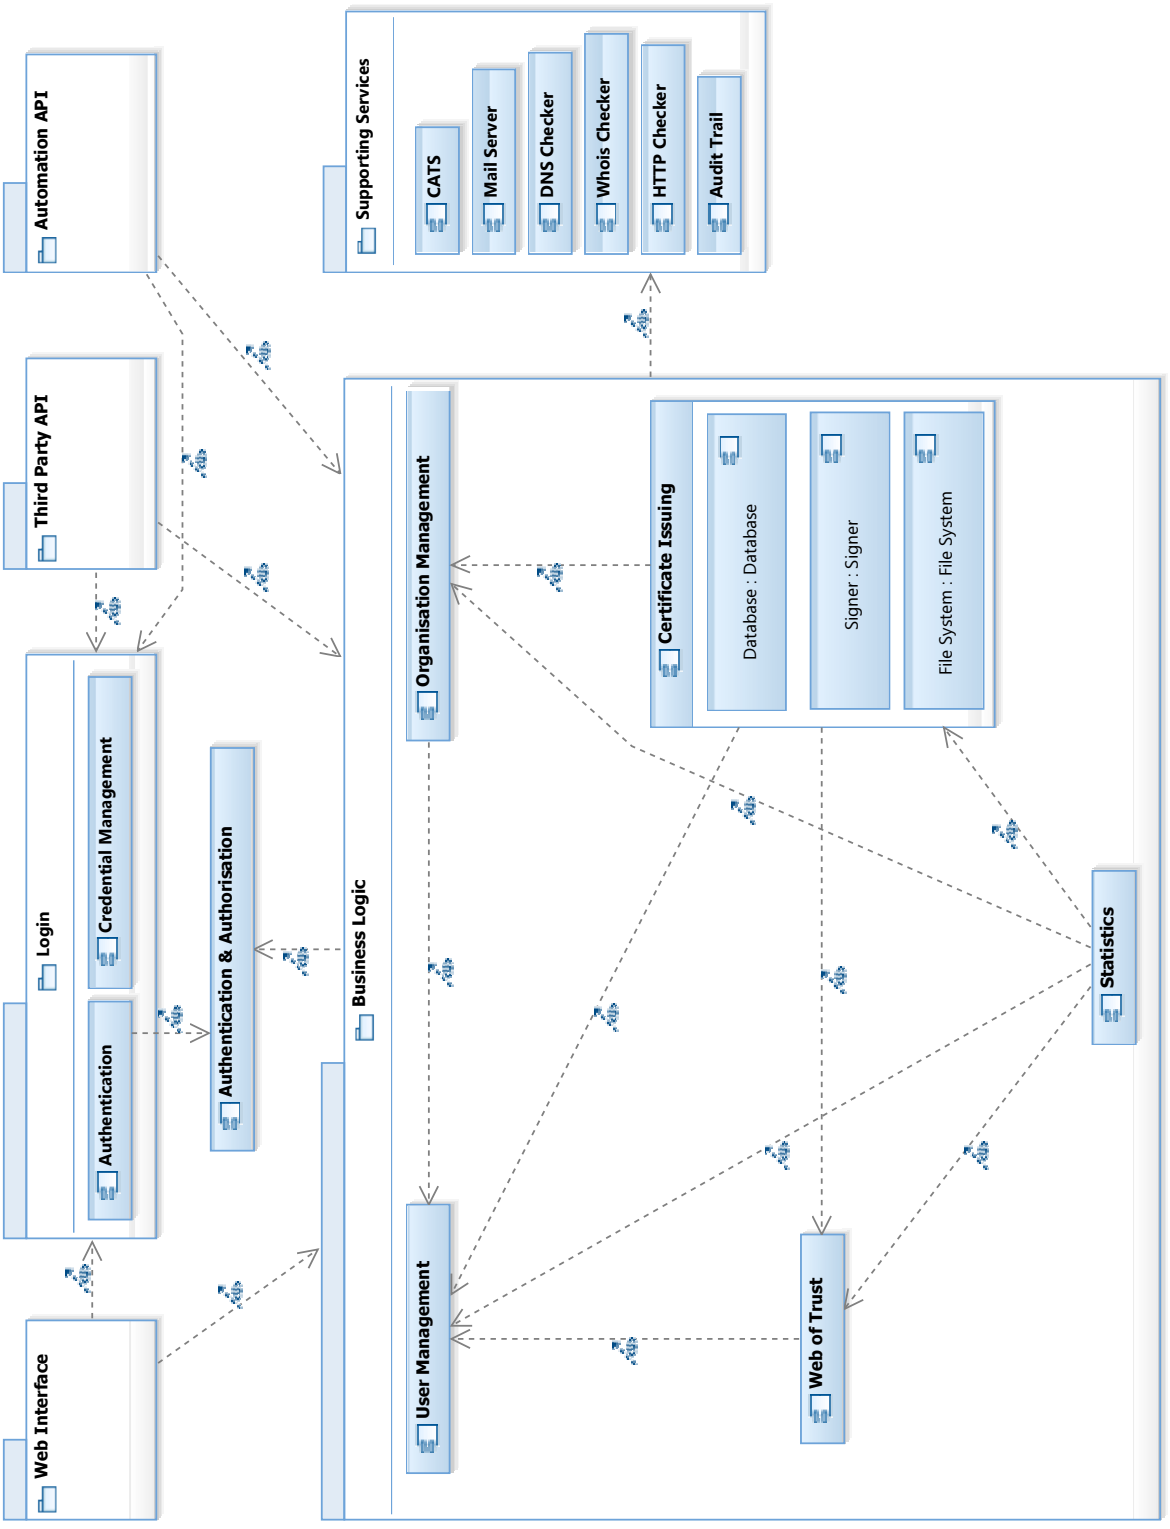

E. Service-Oriented Design Analysis

	Effort/Risk/ Gain
E.1. Web of Trust	
E.1.1. Unauthorised Modification of the Web of Trust	
1. Direct access to the database on the web of trust host	**2/1/8** **5/2/7**
1.1. Recover database credentials AND	**5/1/3**
1.1.1. Recover application credentials	**5/1/3**
1.1.1.1. Exploit information leak (e.g. via error messages, stack traces and other debugging information)	**5/2/2**
1.1.1.2. Exploit memory access violation (buffer overflow attack, printf attack etc.)	**6/1/8**
1.1.1.3. File content leak of configuration files (user controlled paths, unintentional check-ins into version control etc.)	**5/1/3**
1.1.1.4. Set new credentials for application (e.g. via SQL injection)	**6/2/7**
1.1.1.5. Social Engineering	**5/6/6**
1.1.2. Recover administrative credentials	**6/7/8**
1.1.2.1. Social Engineering	**6/7/8**
1.1.2.2. Attack that gains root privileges on the system running the database (exploit of security vulnerability for the OS or other component running with root privileges, access to hardware etc.)	**7/2/8**
1.2. Connect to the database	**4/2/3**
1.2.1. Attack that gains user privileges on the web of trust host	**4/2/5**
2. Inject malicious database commands into the commands sent by the web of trust component (SQL injection)	**6/2/7**
2.1. Bypass input validations (e.g. by using unusual escape sequences that will be unescaped at components closer to the web of trust component) AND	**3/2/3**
2.2. Execute the injected data in the context of SQL commands (e.g. improperly handled variable parts in SQL templates)	**6/1/4**
3. Modify data in the web of trust component	**2/1/7**
3.1. Make a web of trust operation the attacker is authorised to execute, modify the data in unintended ways	**4/1/4**
3.1.1. Exploit operation logic (e.g. sloppy condition checking, buffer overflows)	**4/1/3**
3.2. Execute an operation the attacker is not allowed to execute	**2/1/6**
3.2.1. Bypass authorisation checking in the web of trust component	**6/1/5**

3.2.2. Make the authorisation component report that the attacker is allowed to execute
7/3/7 the operation

7/3/7 3.2.2.1. Directly exploit the authorisation or authentication component

3.2.2.2. Man-in-the-middle attack on the connection to the authorisation component
8/2/5

4/2/5 3.2.2.2.1. Intercept connection to the authorisation component AND

6/4/5 3.2.2.2.1.1. Punch hole through firewall

4/2/5 3.2.2.2.1.2. Attack that gains user privileges on hosts on the internal network

8/1/9 3.2.2.2.2. Attack authentication/encryption protocol between the two components

2/1/6 3.2.3. Use the authentication of another user with the required privileges

2/3/4 3.2.3.1. Social engineering/cooperation by the user (e.g. phishing)

2/1/6 3.2.3.2. Learn authentication credentials of the user (see section E.2.2)

2/1/6 3.2.3.3. Set new authentication credentials for the user (see section E.2.1)

4/2/5 3.2.3.4. Hijack authentication state of the user

3.2.3.4.1. Learn reauthentication credentials (e.g. cookies – analogous to E.2.2) AND
2/1/4

4/2/4 3.2.3.4.2. Pass reauthentication checks, such as IP address restrictions

3.2.3.5. Attack the authentication or reauthentication method (e.g. find a way to
6/1/7 successfully authenticate without knowing the password)

4. Modify the data in transit to the front end or another component (e.g. the certificate
8/2/5 issuing component, see subtree 3.2.2.2.)

E.1.2. Violation of the Confidentiality of the Web of Trust

2/1/5
5/2/7 1. Direct access to database on the web of trust host (see section E.1.1 subtree 1.)

6/2/6 2. Read data from the audit log

6/2/6 2.1. Get read access on the internal network

4/2/3 2.1.1. Connect to the audit log service AND

6/4/5 2.1.1.1. Punch hole through firewall

4/2/5 2.1.1.2. Attack that gains user privileges on hosts on the internal network

6/1/3 2.1.2. Exploit audit log service to read from the log instead of appending to it

6/2/6 2.2. Get read access on the audit trail host

2.2.1. Get administrative access on the audit trail host (see section E.1.1 subtree
7/2/8 1.1.2.2.)

6/2/6 2.2.2. Read audit log with user access on the audit trail host

4/2/3 2.2.2.1. Attack that gains user privileges on the audit trail host AND

2.2.2.2. Attack that gains read privileges on the audit log file to non-administrative users on the audit trail host (via file content leak or user controlled paths in
6/1/4 programmes running as super user etc.)

3. Inject malicious database commands into the commands sent by the web of trust component (SQL injection, see section E.1.1 subtree 2.) **6/2/7**

4. Access data in the web of trust component **2/1/6**

4.1. Make a web of trust operation the attacker is authorised to execute, unintentionally disclose data (see section E.1.1 subtree 3.1.) **4/1/4**

4.2. Execute an operation the attacker is not allowed to execute (see section E.1.1 subtree 3.2.) **2/1/6**

E.2. Login Credentials

E.2.1. Unauthorised Modification of the Login Credentials

2/1/6

1. Direct access to the database on the authentication host (see section E.1.1 subtree 1.) **5/2/7**

2. Inject malicious database commands into the commands sent by the authentication component (SQL injection, see section E.1.1 subtree 2.) **6/2/7**

3. Modify credentials in the authentication component **2/1/6**

3.1. Make an operation of the authentication component the attacker is authorised to execute, modify the data in unintended ways (see section E.1.1 subtree 3.1.) **4/1/4**

3.2. Execute an operation the attacker is not allowed to execute (e.g. set new login credentials) **2/1/6**

3.2.1. Bypass authorisation checking in the authentication component **6/1/5**

3.2.2. Make the authorisation component report that the attacker is authorised to execute the operation (see section E.1.1 subtree 3.2.2.) **7/3/7**

3.2.3. Use the authentication of another user with the required privileges **2/1/6**

3.2.3.1. Social engineering/cooperation by the user (e.g. phishing) **2/3/4**

3.2.3.2. Learn (old) authentication credentials of the user (see section E.2.2) **2/1/6**

3.2.3.3. Pass account recovery methods (e.g. intercept confirmation email sent to the primary email address and learn enough about the user to answer the reset questions) **3/2/6**

3.2.3.4. Hijack authentication state of the user (see section E.1.1 subtree 3.2.3.4.) **4/2/5**

3.2.3.5. Attack the authentication or reauthentication method (e.g. find a way to successfully authenticate without knowing the password) **6/1/7**

E.2.2. Violation of the Confidentiality of the Login Credentials

2/1/7

1. Intercept credentials on the machine of the user (trojan, key logger, cross-site-scripting etc.) **3/2/6**

2. Deceive the user to believe a server controlled by the attacker is the server of CAcert **2/1/6**

2.1. Make the user connect to a server the attacker controls, instead of the real server, and relay the authentication protocol to the real server (phishing, DNS attacks etc.) AND **1/1/4**

2.2. Bypass server side authentication of the encryption protocol (slightly different spelling in the domain name, attacks against domain name comparison like null characters, **2/1/6** use no authentication at all etc.)

3. Intercept credentials on the wire (not possible for challenge/response authentication **8/2/6** protocols)

2/2/3 3.1. Intercept the connection between the user and the server AND

8/1/8 3.2. Break confidentiality of encrypted connection

8/2/6 4. Relay attack on challenge/response protocol

4.1. Route connection through a host the attacker controls (e.g. by attacking name res- **2/2/4** olution, routing protocols or social engineering techniques)

8/1/9 4.2. Break integrity and confidentiality of encrypted connection

5/1/7 5. Intercept credentials on the server

5.1. Attack that gains administrative privileges on the login server (similar to section **6/2/8** E.1.1 subtree 1.1.2.2.)

5.2. Exploit the login server to store credentials and expose them on successive requests **5/1/7**

5.2.1. Persist credentials (e.g. provoke an error that puts the credential as debug infor- **3/1/5** mation in a log file) AND

5/1/6 5.2.2. Restore persisted credentials (e.g. file content leak on a log file)

5.3. Exploit the login server to communicate login credentials to the attacker (e.g. via **5/1/7** code execution attack)

E.3. User Data

E.3.1. Unauthorized Modification of User Data

2/1/8 Analogous to the modification of the web of trust data (section E.1.1).

E.3.2. Unauthorized Access to User Data

2/1/5 Analogous to exposing data from the web of trust (section E.1.2).

E.4. Issued Certificates

E.4.1. Modification of Issued Certificates

5/1/7

1. Direct access to the database of the certificate issuing component with modify access **6/7/8** to the certificate information

1.1. Recover administrative database credentials (see section E.1.1 subtree 1.1.2.) AND **6/7/8**

4/2/3 1.2. Connect to the database (see section E.1.1 subtree 1.2.)

2. Make the certificate issuing component report false meta data (e.g. the user account **5/1/6** the certificate is associated with or whether it may be used for login)

2.1. Exploit the certificate issuing component to report the false meta data (memory **5/1/6** violation, sloppy error checking etc.)

2.2. Modify the data in transit (e.g. to the authentication component, see section E.1.1 **8/2/5** subtree 3.2.2.2.)

E.4.2. Unauthorized Access to Issued Certificates

Analogous to exposing data from the web of trust (section E.1.2). **2/1/5**

E.5. Revocation Information

E.5.1. Unauthorized Modification of Revocation Information

6/7/6

1. Direct access to the database of the certificate issuing component with modify access to the certificate revocation information **6/7/8**

1.1. Recover administrative database credentials (see section E.1.1 subtree 1.1.2.) AND **6/7/8**

1.2. Connect to the database (see section E.1.1 subtree 1.2.) **4/2/3**

2. Modify the data at the point of distribution into the system of OCSP responders and CRL caches **8/2/6**

2.1. Forge the signature on the OCSP response or CRL AND **7/2/6**

2.1.1. Get access to the certificate signing keys (see section E.7.2) **7/2/8**

2.1.2. Forge the cryptographic signature **9/1/9**

2.2. Modify the data on the wire (man in the middle attack) **8/2/7**

2.2.1. Route the connection from the OCSP responder or CRL cache to the authoritative revocation server to a host the attacker controls (e.g. by attacking name resolution, routing protocols or social engineering techniques) AND **2/2/4**

2.2.2. Break the authenticity of the encrypted connection **8/1/8**

3. Modify the data in the system of OCSP responders and CRL caches **7/2/6**

3.1. Forge the signature on the OCSP response or CRL (see subtree 2.1.) AND **7/2/6**

3.2. Get write access to an OCSP responder or CRL cache **4/2/6**

4. Modify the data on the wire between the client and the system of OCSP responders and CRL caches **7/2/6**

4.1. Forge the signature on the OCSP response or CRL (see subtree 2.1.) AND **7/2/6**

4.2. Route the connection from the client to an OCSP responder or CRL cache to a host the attacker controls (e.g. by attacking name resolution, routing protocols or social engineering techniques) **2/2/6**

E.5.2. Prevent Access to Revocation Information

5/3/5

1. Route a significant number of connections from clients to OCSP servers or CRL caches to a host the attacker controls or a dead end (e.g. by attacking name resolution, routing protocols or social engineering techniques) **5/5/6**

2. Denial of service attack against each of the distributed OCSP responders and CRL caches **6/3/5**

3. Denial of service attack against the point of distribution into the system of OCSP responders and CRL caches (single point of failure but possible to use white list filtering in the router) **5/3/3**

4. Denial of service attack against the signing server or certificate issuing component (e.g. by requesting and revoking many certificates at once) **3/8/7**

E.6. Root/Subroot Certificates

E.6.1. Modification of Root/Subroot Certificates

6/2/6
6/2/8 1. Write access to the file system on the front end server

1.1. Attack that gains root privileges on the front end server (similar to section E.1.1
6/2/8 subtree 1.1.2.2.)

1.2. Exploit some component or service running with elevated privileges on the front end
6/2/8 server to overwrite files

E.7. Certificate Signing Keys

E.7.1. Modification of the Certificate Signing Keys

8/2/6
1. Exploit some component or service running with elevated privileges on the signing
8/2/6 server to overwrite the certificate signing keys AND

6/2/8 2. Access the serial connection (see section E.7.2 subtree 1.2.)

E.7.2. Violation of the Confidentiality of the Certificate Signing Keys

7/2/8
7/2/8 1. Exploit the signing server via the serial access protocol

1.1. Exploit the signing server to expose information about the key material (e.g. by causing part of the signing key material to be included in a signed certificate via
7/2/8 memory violation) AND

6/2/8 1.2. Access the serial connection

1.2.1. Directly access the serial connection (e.g. by getting root access on the host
7/2/8 running the certificate issuing component, see section E.1.1 subtree 1.1.2.2.)

1.2.2. Route malicious data through the normal signing process (similar to the SQL
6/2/7 injection, see section E.1.1 subtree 2.)

7/5/7 2. Gather information about the signing key material by observing side channels

2.1. Discover a side channel (timings, power consumption, cache access patterns etc.)
5/2/4 AND

2.2. Access the side channel (e.g. by analysing response times or measuring power con-
7/5/4 sumption)